Supplement of

From Points to Images: deep-learning enhanced spatial-temporal rainfall modeling from point measurements

Bing-Zhang Wang, Li-Pen Wang, Auguste Gires

Correspondence to: Li-Pen Wang (lpwang@ntu.edu.tw)

S0 Introduction

This supplement is organised into two main sections. Section S1 details the implementation of the DeepKriging and Spatio-temporal DeepKriging baselines used in this study, including the modifications made to the original positional encoding, the direct conditioning on rainfall observations, and the revised training loss used for fair comparison with the proposed framework. Section S2 provides example snapshot images for all selected storm events for the Stage I and the Stage II experiments. This section is intended to support the quantitative results in the main manuscript by illustrating how the compared methods behave across different rainfall events.

S1 DeepKriging and Spatio-temporal DeepKriging baseline implementation

S1.1 Original DeepKriging and Spatio-temporal DeepKriging and their applications

The DeepKriging (DK) framework was originally proposed as a deep-learning-based spatial interpolation approach and evaluated on large-scale environmental datasets over the United States, including air quality and climate-related variables exhibiting complex spatial dependence (Chen et al., 2022). The underlying idea of DK is to reformulate spatial interpolation as a regression problem, in which spatial dependence is represented through a set of basis functions rather than an explicitly specified covariance model. Specifically, the input coordinates are first transformed through a multi-resolution positional encoding, enabling DK to capture spatial variability across multiple scales. These encoded representations are then fed into a multilayer perceptron (MLP) that learns the mapping between location and the corresponding variable of interest. The formulation is flexible and, in principle, allows other predictors (e.g., rainfall intensity or other meteorological variables) to be incorporated alongside spatial coordinates.

Spatio-temporal DeepKriging (STDK) is an extension of DK. It introduces a temporal positional encoding, representing both spatial and temporal dependencies. In the proposed method, the encoded representations were coupled with ConvLSTM-based forecasting components, demonstrating the applicability of this approach to spatio-temporal rainfall reconstruction and prediction tasks (Nag et al., 2023). These characteristics make DK and STDK suitable deep-learning baselines for spatio-temporal rainfall interpolation problems.

However, several practical limitations arise when applying DK-based formulations within the present study. First, both DK and STDK are fundamentally designed for point-wise prediction, which leads to high computational cost when producing dense gridded rainfall fields. Second, the original formulations primarily rely on the positional encoding as input, without explicitly incorporating observed rainfall measurements as additional predictors. Following the recommendations given in the original literature, observed gauge values are concatenated in this study with the positional encoding vectors before being passed to the MLP. This allows the network to directly utilise available rainfall information while preserving the original architecture. Third, the original DK training strategy evaluates loss only on validation locations, differing from the experimental design adopted in this study. To ensure consistency with the P2I-GAN framework and enable fair comparison, model performance is evaluated over all non-training locations, and the original loss formulation is replaced with the weighted objective used in this study (weighted L1, KL divergence, and pm_loss). A comparison between the original DK loss and the modified objective function is also provided next in this Supplement. The implementation details corresponding to these modifications are described in the following sections.

S1.2 Computational bottleneck of dense positional encoding

DeepKriging represents spatial dependence by augmenting the network input with multi-resolution radial basis functions rather than using raw coordinates alone. In the original formulation, the spatial embedding is constructed from compactly supported Wendland basis functions defined on a rectangular grid of knots. As the spatial resolution increases, the number of basis functions grows rapidly in two-dimensional applications. The original DeepKriging paper gives a four-level two-dimensional example with

$$10^2 + 19^2 + 37^2 + 73^2 = 7159$$

Table S1. Comparison of positional encoding schemes under the present dense-grid setting. Here H, W, and T denote the grid height, grid width, and sequence length, respectively. Memory values refer to positional encoding only under `float32` storage.

Method	Positional encoding	Basis dimension	Positional encoding size
DK (original practical use)	Full 2-D Wendland multi-resolution basis	7159	83.89 GiB
STDK (original practical use)	2-D Wendland + 1-D temporal basis	297	3.48 GiB
DK (modified)	Subsampled 2-D Wendland basis	139	8.69 MiB
STDK (modified)	Subsampled 2-D Wendland + subsampled 1-D basis	$139+44$	8.69 MiB

spatial basis functions, and the default regression network uses three hidden layers with 100 neurons each. In the original STDK paper, spatial dependence is represented by Wendland radial basis functions and temporal dependence by Gaussian radial basis functions. To reduce the computational burden, the original paper avoids the full tensor-product basis of size GH and instead stacks the spatial and temporal encodings into a feature vector of size $G+H$. In the competition setting reported in that paper, the interpolation model used three spatial basis levels with $G=25,81,144$ and three temporal basis levels with $H=10,15,45$, while the paper reports 297 basis functions in the final implementation.

These original DK and STDK formulations were designed for point-wise interpolation, where each queried location is processed individually. This design is manageable in their original applications, where prediction and evaluation are performed on sparse sites or held-out locations. However, the present study targets dense radar-like rainfall fields, where the prediction domain is an entire image grid rather than a sparse set of points. For an input of size $12\times16\times128\times128$ in $BTHW$ format, the total number of queried pixel-time locations is $3{,}145{,}728$. In such a setting, directly extending the original point-wise formulations to all queried locations becomes computationally expensive, because the positional encoding and the shared MLP must be evaluated repeatedly for every pixel in every frame.

To make DK-type baselines feasible under this dense-grid setting, we preserve the original multi-resolution Wendland formulation but simplify the knot construction through deterministic subsampling. Instead of constructing the full set of knots at each level, we retain a fixed number of representative knots per level, resulting in a compact low-rank spatial embedding. In the modified DK implementation, the spatial basis is reduced to

$$\tilde{K}_s = 10+19+37+73 = 139.$$

A similar simplification is applied in the temporal branch of the modified STDK implementation. For $T=16$, the modified temporal encoder yields $\tilde{K}_t=44$, so that the positional encoding dimension in the modified STDK implementation becomes

$$\tilde{K}_s+\tilde{K}_t = 183.$$

Formally, the original dense spatial embedding can be written as

$$\phi_s(\mathbf{u}) = \left[\phi_1(\mathbf{u}),\dots,\phi_{K_s}(\mathbf{u})\right]^\top,$$

where K_s denotes the number of spatial knots used in the original formulation. In the modified implementation, this is replaced by

$$\tilde{\phi}_s(\mathbf{u}) = \left[\phi_1(\mathbf{u};\tilde{\mathbf{u}}_1,\tilde{\theta}_1),\dots,\phi_{\tilde{K}_s}(\mathbf{u};\tilde{\mathbf{u}}_{\tilde{K}_s},\tilde{\theta}_{\tilde{K}_s})\right]^\top, \qquad \tilde{K}_s = 139,$$

where $\tilde{\mathbf{u}}_j$ are the retained knot locations after subsampling. Since subsampling increases the effective spacing between retained knots, the support radius is slightly enlarged in the modified implementation to preserve overlap between neighboring basis functions and suppress localized knot-imprinting artifacts.

For the modified STDK implementation, the input to the shared MLP becomes

$$\tilde{\mathbf{x}}(\mathbf{u},t) = \left[\tilde{\phi}_s(\mathbf{u}),\tilde{\phi}_t(t),\mathbf{z}^{\mathrm{obs}}_{1:T}\right]^\top,$$

where $\mathbf{z}_{1:T}^{\mathrm{obs}} \in \mathbb{R}^{79T}$ denotes the sequence of observed rainfall values concatenated over time. Compared with the original formulation, the main computational reduction therefore comes from shrinking the positional encoding itself before the MLP stage, rather than from changing the shared prediction head.

To illustrate the practical difference under the present dense-grid setting, we estimate the size of the positional encoding if the original point-wise basis dimensions are naively expanded to all queried pixel-time locations. All calculations below assume single-precision floating point storage (`float32`), so each value occupies 4 bytes.

Under this setting, extending the original DeepKriging example with 7159 basis functions to all queried locations would require about 83.89 GiB for the positional encoding alone. Using the 297 basis functions reported in the original STDK implementation would still require about 3.48 GiB. These estimates do not include observed inputs, hidden activations, gradients, or optimizer states, and therefore reflect only the positional component.

In contrast, the modified DK implementation used in this study does not materialize a full $B \times T \times H \times W \times \tilde{K}_s$ positional tensor. Instead, it caches the spatial basis only once at image level and reuses it frame by frame. With $H = W = 128$ and $\tilde{K}_s = 139$, this cached basis requires only about 8.69 MiB. Likewise, the modified STDK implementation caches the spatial basis and a small temporal basis, leading to a nearly identical cached positional size of about 8.69 MiB. Relative to the dense-grid expansion of the original practical settings, the cached positional basis is therefore reduced to approximately 0.010% of the corresponding DK scale and 0.244% of the corresponding STDK scale.

Overall, the original DK and STDK papers were designed for point-wise interpolation, where the positional basis is evaluated per queried location rather than over a dense image grid. When these same practical basis settings are extended to the dense-grid rainfall interpolation problem considered here, the positional representation becomes substantially more expensive. The proposed subsampling strategy and implementation changes therefore make DK and STDK computationally feasible in the present setting while preserving the main positional encoding logic of the original models.

S1.3 Direct conditioning on observational data

Another important difference between the original DK/STDK formulations and the present study lies in how observational information is incorporated into the model input. In the original DeepKriging formulation, the network input consists of the positional embedding together with external covariates. Likewise, in the original STDK interpolation model, the input is defined as the concatenation of the spatio-temporal basis vector and additional covariates. In both cases, the model is not designed to directly ingest the full set of observed target values as predictors.

This differs from the rainfall interpolation setting considered in this study. Here, the available information at each time step is a sparse set of gauge observations, whereas the prediction target is a complete gridded rainfall field. To make DK-type models compatible with this setting, the observed rainfall values are directly concatenated to the positional encoding before entering the MLP. This preserves the original feed-forward prediction logic, but changes the conditioning information from external environmental covariates to the actually observed rainfall measurements.

In the modified DK implementation, the model first computes the spatial basis for all queried pixels, then selects 79 visible observations from the current frame according to the observation mask, and concatenates these 79 observed values to the spatial basis at every queried pixel. The prediction for each frame is therefore conditioned on the same frame-wise observation vector. The resulting input representation can be written as

$$\mathbf{x}^{\mathrm{DK}}(\mathbf{u}, t) = \left[\tilde{\phi}_s(\mathbf{u}), \mathbf{z}_t^{\mathrm{obs}}\right]^\top,$$

Table S2. Pseudo-code summary of how observational rainfall values are incorporated in the modified DK and STDK baselines.

DK (modified)	STDK (modified)
1. Compute $\tilde{\phi}_s(\mathbf{u})$ for all pixels.	1. Compute $\tilde{\phi}_s(\mathbf{u})$ and $\tilde{\phi}_t(t)$ for all pixel-time pairs.
2. Select 79 visible observations from the current frame.	2. Select 79 visible observations from each frame.
3. Form $\mathbf{z}_t^{\mathrm{obs}} \in \mathbb{R}^{79}$.	3. Concatenate them into $\mathbf{z}_{1:T}^{\mathrm{obs}} \in \mathbb{R}^{79T}$.
4. For each queried pixel, use $[\tilde{\phi}_s(\mathbf{u}), \mathbf{z}_t^{\mathrm{obs}}]$.	4. For each queried pixel-time pair, use $[\tilde{\phi}_s(\mathbf{u}), \tilde{\phi}_t(t), \mathbf{z}_{1:T}^{\mathrm{obs}}]$.
5. Feed the concatenated vector to the shared MLP.	5. Feed the concatenated vector to the shared MLP.

Table S3. Categorical verification scores for DK and STDK under different training loss settings across four rainfall thresholds (0.5, 2.0, 4.0, 8.0 mm/h). Here, *Ours* denotes the weighted objective used in this study, and *Ours w/o pm_loss* denotes the same objective without pm_loss.

MODELS	$\geq$ 0.5 mm/h				$\geq$ 2.0 mm/h				$\geq$ 4.0 mm/h				$\geq$ 8.0 mm/h			
	POD	FAR	CSI	HSS	POD	FAR	CSI	HSS	POD	FAR	CSI	HSS	POD	FAR	CSI	HSS
DK (MSE)	0.742	0.148	0.657	0.737	0.377	0.234	0.338	0.488	0.139	0.280	0.132	0.231	0.017	0.415	0.017	0.033
STDK (MSE)	0.714	0.149	0.635	0.718	0.267	**0.195**	0.251	0.385	0.043	**0.248**	0.043	0.081	0.000	**0.000**	0.000	0.000
DK (Ours)	0.823	0.257	0.640	0.712	**0.559**	0.460	0.379	0.524	**0.318**	0.711	0.178	0.296	**0.125**	0.884	**0.064**	**0.119**
STDK (Ours)	0.826	0.296	0.613	0.682	0.441	0.465	0.319	0.458	0.122	0.770	0.086	0.154	0.014	0.965	0.010	0.020
DK (w/o L_{pm})	0.846	0.226	**0.679**	**0.749**	0.537	0.313	**0.432**	**0.584**	0.282	0.426	**0.233**	**0.374**	0.066	0.564	0.061	0.114
STDK (w/o L_{pm})	**0.851**	0.299	0.625	0.692	0.461	0.415	0.347	0.493	0.199	0.567	0.158	0.268	0.006	0.924	0.006	0.011

where $\mathbf{z}_t^{\mathrm{obs}} \in \mathbb{R}^{79}$ denotes the visible gauge values at time step t.

In the modified STDK implementation, this idea is extended across time. Spatial and temporal basis functions are first computed for all queried pixel-time pairs. The model then selects 79 visible observations from each frame, concatenates them over time, and appends the resulting sequence-wide observation vector to every queried pixel-time pair. The input representation is therefore

$$\mathbf{x}^{\mathrm{STDK}}(\mathbf{u}, t) = \left[\tilde{\phi}_s(\mathbf{u}), \tilde{\phi}_t(t), \mathbf{z}_{1:T}^{\mathrm{obs}}\right]^\top,$$

where $\mathbf{z}_{1:T}^{\mathrm{obs}} \in \mathbb{R}^{79T}$ contains the observed rainfall values from all input frames.

This modification is especially important because the original DK and STDK studies were developed for $\mathrm{PM}_{2.5}$ applications with rich covariate fields, whereas the present problem is driven primarily by sparse rainfall observations. Direct concatenation of gauge values therefore provides the simplest way to retain the original DK/STDK architecture while allowing the model to condition on the physically relevant information in the rainfall interpolation task.

S1.4 Differences in training loss formulation

The original DeepKriging paper adopts mean squared error (MSE) as the default loss for regression tasks, while the original STDK interpolation framework is likewise trained under a standard regression objective. This choice is suitable for the original $\mathrm{PM}_{2.5}$ applications, where the target variable is continuous and does not exhibit the same degree of intermittency and skewness as rainfall. However, rainfall fields are highly zero-inflated, strongly skewed, and dominated by rare but meteorologically important extremes. Under such a setting, directly applying MSE loss tends to emphasize average reconstruction accuracy while underweighting the categorical and heavy-rainfall characteristics that are more relevant in precipitation evaluation.

In practice, training DK and STDK with the original MSE objective leads to noticeably poorer categorical performance at moderate and heavy rainfall thresholds, especially beyond $4\ \mathrm{mm\ h^{-1}}$ (Table S3). This does not necessarily imply that the architectures themselves are unable to learn useful rainfall structures. Rather, it suggests that the original loss formulation does not sufficiently reflect the distributional characteristics of rainfall intensity. To isolate this effect, we compare the original MSE-trained versions with versions retrained using the same weighted objective adopted in this study.

Specifically, the modified DK and STDK baselines are optimized using a weighted combination of weighted L_1 loss, KL divergence, and pm_loss. This change serves two purposes. First, it aligns the optimization target across all neural baselines, so that the formal comparison focuses on architectural differences rather than on inconsistent training objectives. Second, it improves the sensitivity of DK-type models to moderate and high rainfall intensities that would otherwise be underrepresented under MSE-based optimization. Quantitatively, the use of the weighted objective leads to a clear improvement in categorical rainfall skill relative to the original MSE-based setting, as summarized in Table S3.

At the same time, the visual results reveal an additional limitation of DK-type models under the modified loss. As shown in Fig. S1, when $L_{\mathrm{pm}}$is included, both DK and STDK frequently generate spatially unnatural high-value patches. This behavior arises because L_{pm}explicitly emphasizes the upper tail of the rainfall distribution, while the relatively simple basis-plus-MLP architecture lacks the flexibility to reproduce diverse extreme-rainfall morphologies. As a result, the model may satisfy the upper-tail constraint by introducing a compact rainfall region in locations where it minimally affects the remaining loss terms.

This artifact largely disappears once L_{pm} is removed, as can also be seen in Fig. S1, indicating that the issue is induced by the interaction between extreme-value supervision and limited model expressiveness, rather than by positional encoding alone. Nevertheless, for fairness, the main study retains the full weighted objective for all neural models, including DK and STDK, and reports the MSE-based versions only in the Supplement. This ensures that the final comparison reflects differences in model architecture under a shared optimization target, rather than improvements introduced solely by a more favorable loss design.

S1.5 Summary of baseline modifications

This supplementary section described three practical modifications that enabled DK and STDK to be fairly compared within the framework of this study. The results show that these adjustments substantially reduced the computational burden, making both models feasible for dense-grid rainfall interpolation. In particular, simplifying the positional encoding reduced the effective computational demand to approximately 1% of that required by the original dense formulation, while preserving the core modeling logic.

The experiments also show that, with an appropriate loss design, both DK and STDK can achieve clear performance improvements over their original MSE-based formulations. This indicates that part of their limited performance in rainfall interpolation originates from the mismatch between the original loss function and the statistical characteristics of rainfall, rather than from the positional encoding framework alone.

At the same time, the present results also reveal an architectural limitation of these models. The current MLP backbone uses hidden layers of width 100, while the actual input dimension is already large due to the concatenation of positional encodings and observed gauge values. As a result, potentially useful information may be compressed too aggressively at the first layer and may not be fully utilized by the network. Therefore, improving the feature extraction and representation mechanism within the backbone remains a promising direction for future work.

From a modeling perspective, DK and STDK are particularly meaningful in situations where only point observations are available and no gridded field data exist. Under such conditions, their positional encoding design allows the model to learn dependencies among sparse observation sites and reconstruct a continuous spatial field from point-based information alone. Although this encoding strategy is memory-intensive, and although other forms of spatial information such as terrain features or simpler location descriptors may be learned more efficiently, DK- and STDK-type formulations retain a unique advantage in point-only settings. In this sense, they address a scenario that cannot be directly handled within the P2I-GAN framework.

Overall, without altering the core model definitions, we introduced minimal adjustments to the positional encoding, observational input construction, and evaluation protocol, and trained the models using the same loss function adopted in this study. These modifications made DK and STDK practically applicable in the present experiments and enabled a fairer comparison of the model architectures themselves.

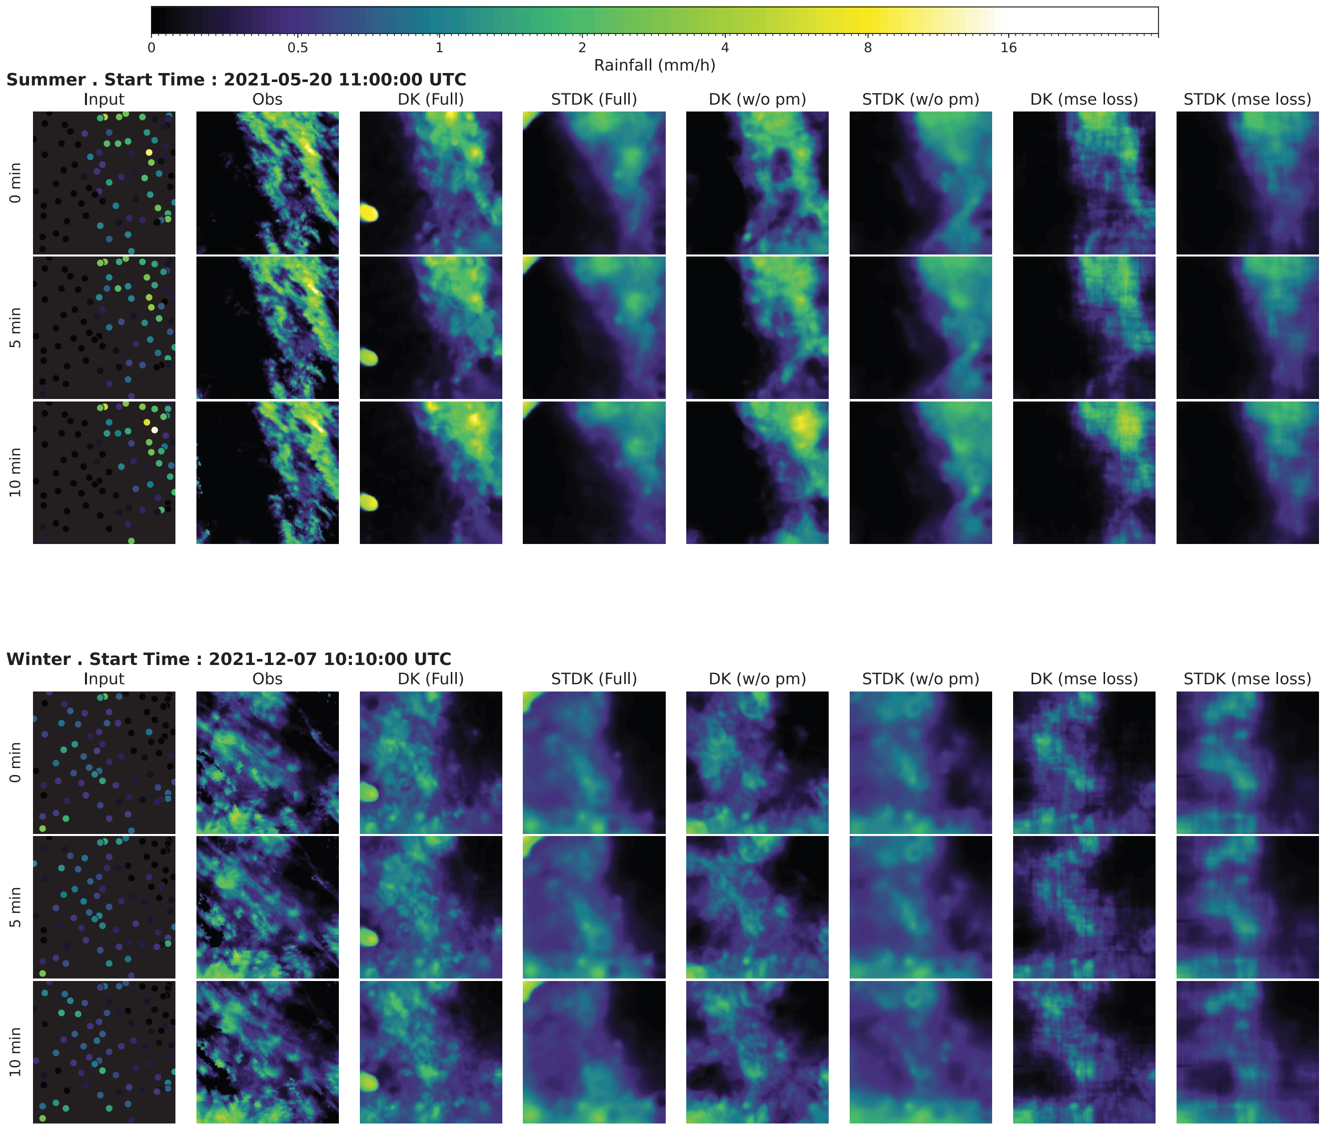

Figure S1. Visual comparison of DK and STDK under different training settings. The figure compares models trained with the original MSE objective and with the weighted objective used in this study. Although the weighted objective improves rainfall intensity reconstruction and categorical skill, it introduces localized unnatural high-value patches in DK and STDK outputs when $L_{\mathrm{pm}}$is included.

S2 Event-based visual verification cases

This section provides additional example snapshots for the storm cases used in the result verification of the main manuscript. A total of 20 representative storm events were selected to cover a range of rainfall durations and accumulation magnitudes, allowing the visual performance of the models to be examined under diverse precipitation conditions. Following the structure in this study, one representative interpolated field is provided for each of the 20 events in Stage I, and one representative gauge-derived reconstruction field is provided for each of the selected 10 events in Stage II.

Table S4 (the same as Table 1 in the main manuscript) summarises the selected events together with their verification statistics derived from the MIDAS rain-gauge product (RG) and the Nimrod radar product (RD). Specifically, the table reports the event start time, end time, duration, and the corresponding maximum and mean rainfall intensities. For the 2021 events (Events 1–10), both RG and RD statistics are available, whereas for the 2022 events (Events 11–20), only RD statistics are available. This table is included to provide the event context for the visual examples presented in the following sections, so that each figure can be interpreted together with the event duration and overall rainfall magnitude.

Table S4. Selected storm events for result verification and summary statistics from MIDAS rain gauges (RG) and Nimrod radar (RD). RG statistics are calculated from 99 gauge locations, whereas RD statistics are calculated over the full radar field (16,384 pixels). For the 2021 events (Events 1–10), both RG and RD statistics are available; for the 2022 events (Events 11–20), only RD statistics are available.

Event	Start time	End time	Duration (hours)	Summary statistics (mm/h) Max. RG	Max. RD	Mean RG	Mean RD
1	2021-01-15 22:00	2021-01-16 18:00	20	16.8	48.0	0.313	0.424
2	2021-03-09 22:00	2021-03-11 00:00	26	17.4	13.8	0.220	0.285
3	2021-05-03 08:00	2021-05-04 05:00	21	47.3	45.8	0.722	0.642
4	2021-05-08 01:00	2021-05-08 21:00	20	13.4	21.2	0.878	0.670
5	2021-05-20 09:00	2021-05-22 02:00	41	21.0	32.7	0.368	0.386
6	2021-05-23 11:00	2021-05-24 07:00	20	34.8	43.6	0.325	0.288
7	2021-07-05 17:00	2021-07-06 13:00	20	22.6	85.4	0.530	0.432
8	2021-10-02 05:00	2021-10-03 02:00	21	13.0	23.4	0.377	0.334
9	2021-10-31 03:00	2021-11-01 07:00	28	89.5	94.0	0.627	0.526
10	2021-12-07 08:00	2021-12-08 05:00	21	39.2	29.7	0.293	0.355
11	2022-01-08 04:00	2022-01-08 22:00	18	N/A	31.2	N/A	0.221
12	2022-03-16 11:00	2022-03-17 05:00	18	N/A	16.7	N/A	0.603
13	2022-04-12 06:00	2022-04-12 19:00	13	N/A	119.4	N/A	0.337
14	2022-05-11 02:00	2022-05-11 20:00	18	N/A	85.4	N/A	0.417
15	2022-05-15 20:00	2022-05-16 11:00	15	N/A	89.6	N/A	0.287
16	2022-06-18 15:00	2022-06-19 06:00	15	N/A	55.4	N/A	0.459
17	2022-08-25 05:00	2022-08-25 19:00	14	N/A	45.8	N/A	0.232
18	2022-09-30 09:00	2022-10-01 03:00	18	N/A	31.2	N/A	0.474
19	2022-10-31 18:00	2022-11-01 12:00	18	N/A	77.6	N/A	0.580
20	2022-11-23 03:00	2022-11-23 19:00	16	N/A	85.4	N/A	0.506

S2.1 Stage I: 20 event cases

The following event-by-event cases provide one representative interpolation figure for each of the 20 selected storms.

Event 1

Event 1 spans from 2021-01-15 22:00 to 2021-01-16 18:00, with a total duration of 20 hours. According to Table S4, the reported maximum rainfall intensities is 48.0 mm/h for the radar product (RD), while the corresponding mean intensities is 0.424 mm/h. Figure S2 presents one representative interpolation case for this event.

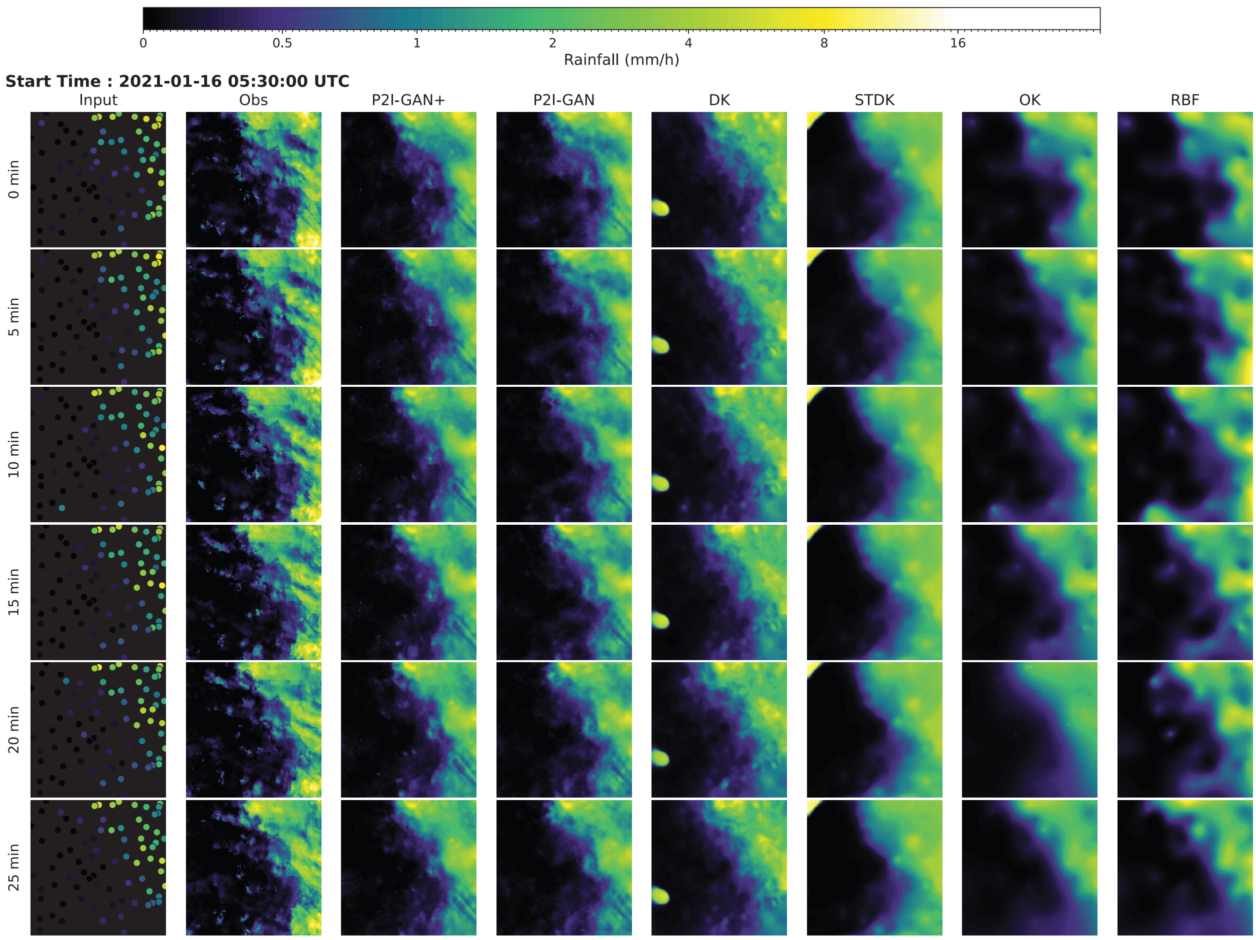

Figure S2. Representative interpolation case for Event 1 (2021-01-15 22:00 to 2021-01-16 18:00).

Event 2

Event 2 spans from 2021-03-09 22:00 to 2021-03-11 00:00, with a total duration of 26 hours. According to Table S4, the reported maximum rainfall intensities is 13.8 mm/h for the radar product (RD), while the corresponding mean intensities is 0.285 mm/h. Figure S3 presents one representative interpolation case for this event.

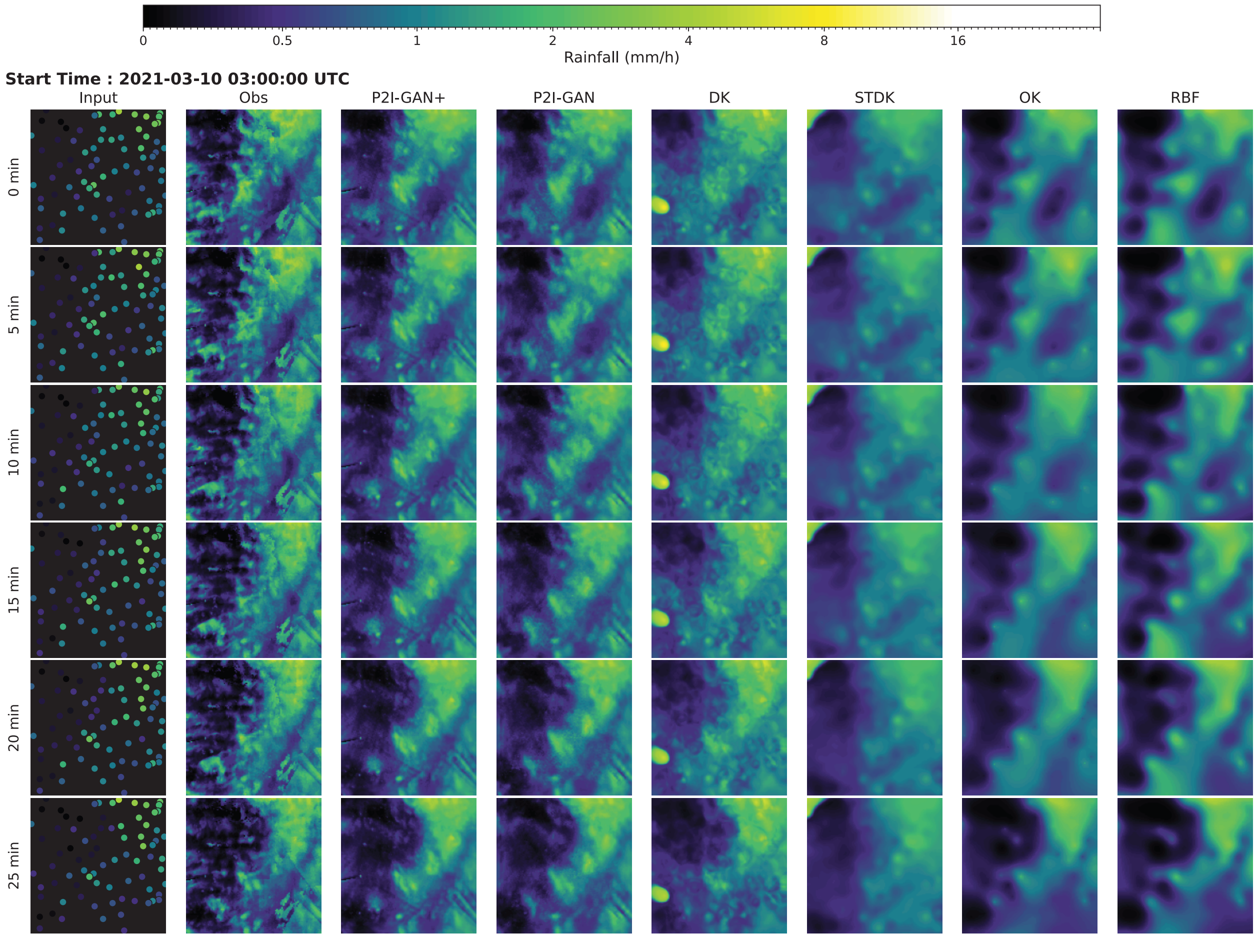

Figure S3. Representative interpolation case for Event 2 (2021-03-09 22:00 to 2021-03-11 00:00).

Event 3

Event 3 spans from 2021-05-03 08:00 to 2021-05-04 05:00, with a total duration of 21 hours. According to Table S4, the reported maximum rainfall intensities is 45.8 mm/h for the radar product (RD), while the corresponding mean intensities is 0.642 mm/h. Figure S4 presents one representative interpolation case for this event.

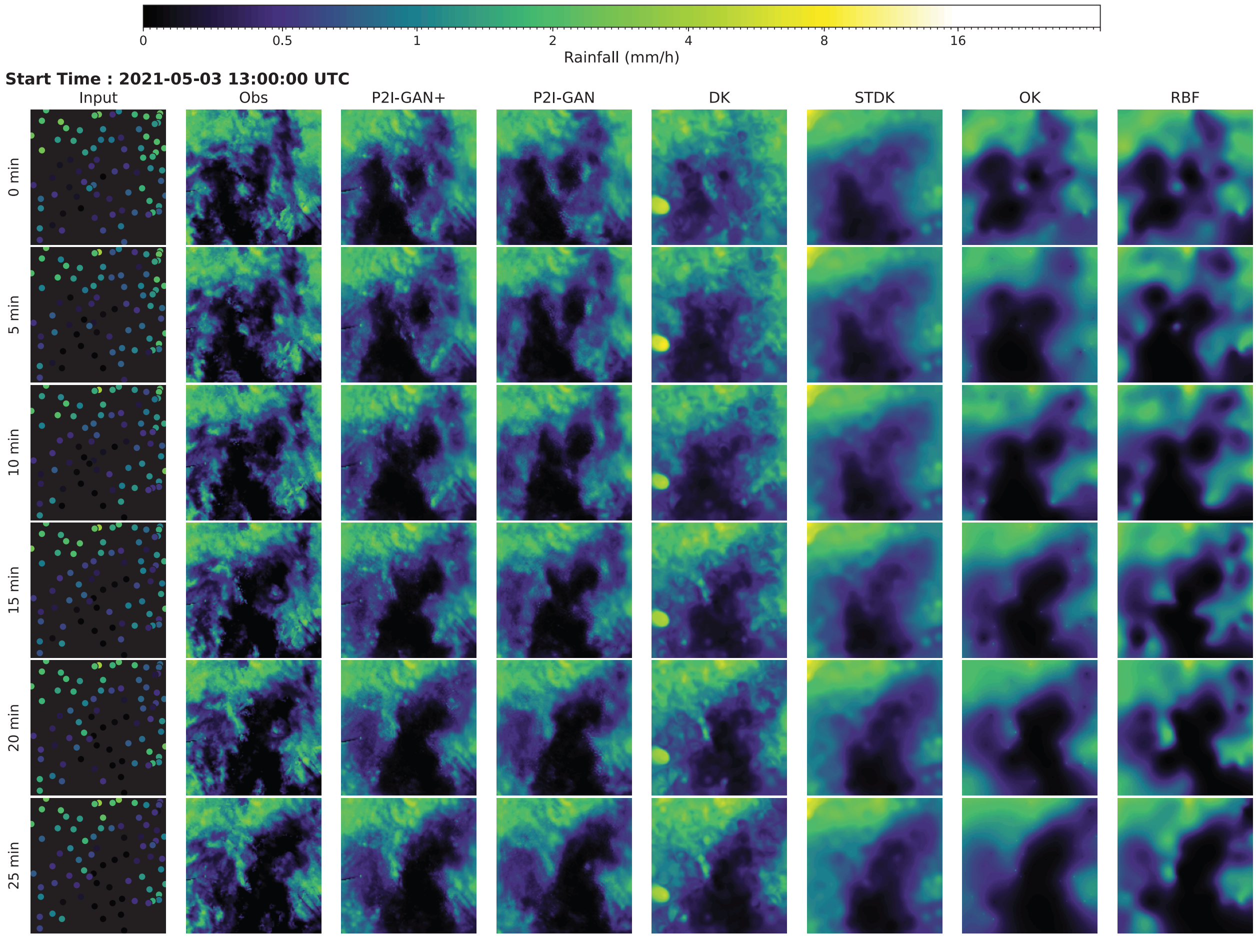

Figure S4. Representative interpolation case for Event 3 (2021-05-03 08:00 to 2021-05-04 05:00).

Event 4

Event 4 spans from 2021-05-08 01:00 to 2021-05-08 21:00, with a total duration of 20 hours. According to Table S4, the reported maximum rainfall intensities is 21.2 mm/h for the radar product (RD), while the corresponding mean intensities is 0.670 mm/h. Figure S5 presents one representative interpolation case for this event.

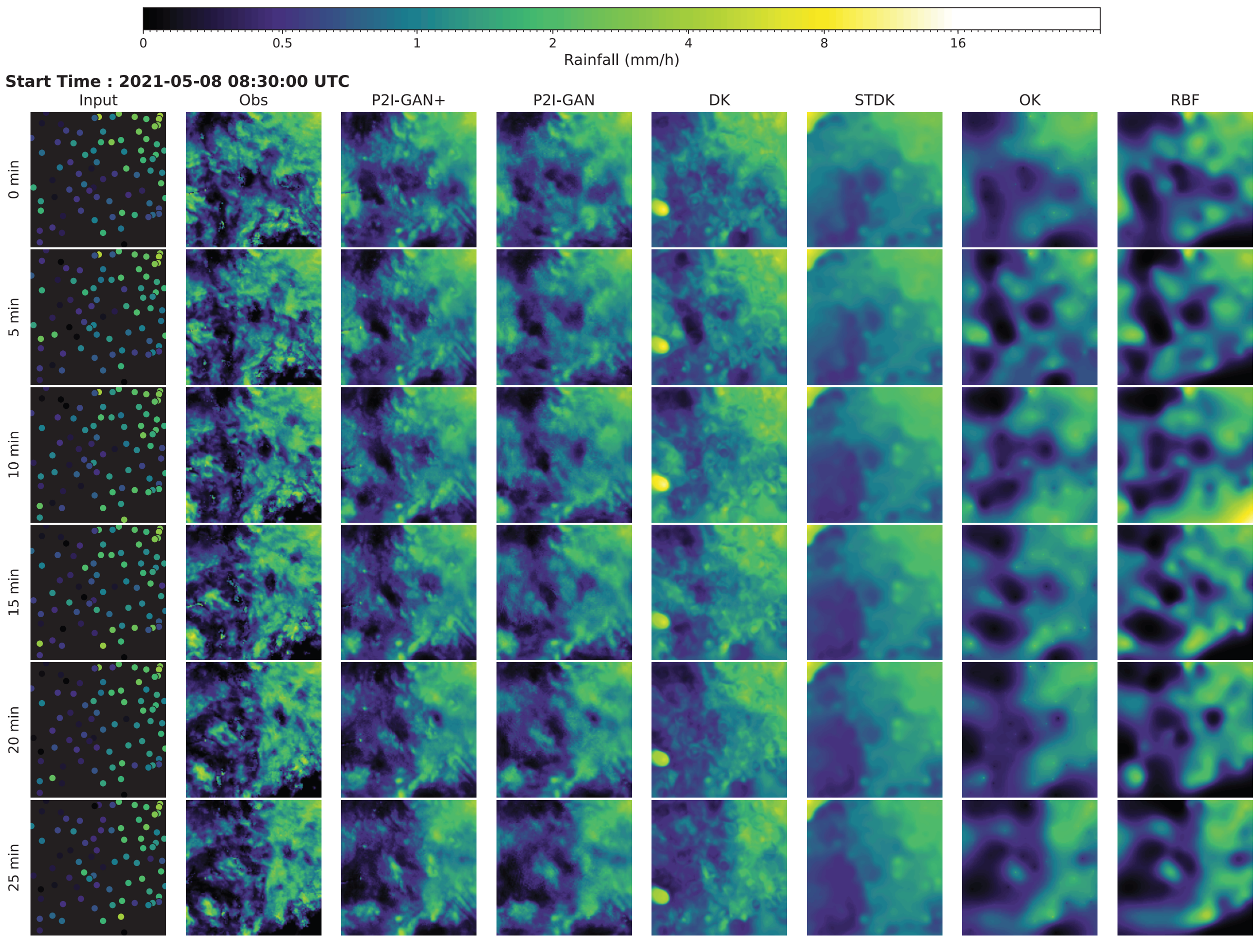

Figure S5. Representative interpolation case for Event 4 (2021-05-08 01:00 to 2021-05-08 21:00).

Event 5

Event 5 spans from 2021-05-20 09:00 to 2021-05-22 02:00, with a total duration of 41 hours. According to Table S4, the reported maximum rainfall intensities is 32.7 mm/h for the radar product (RD), while the corresponding mean intensities is 0.386 mm/h. Figure S6 presents one representative interpolation case for this event.

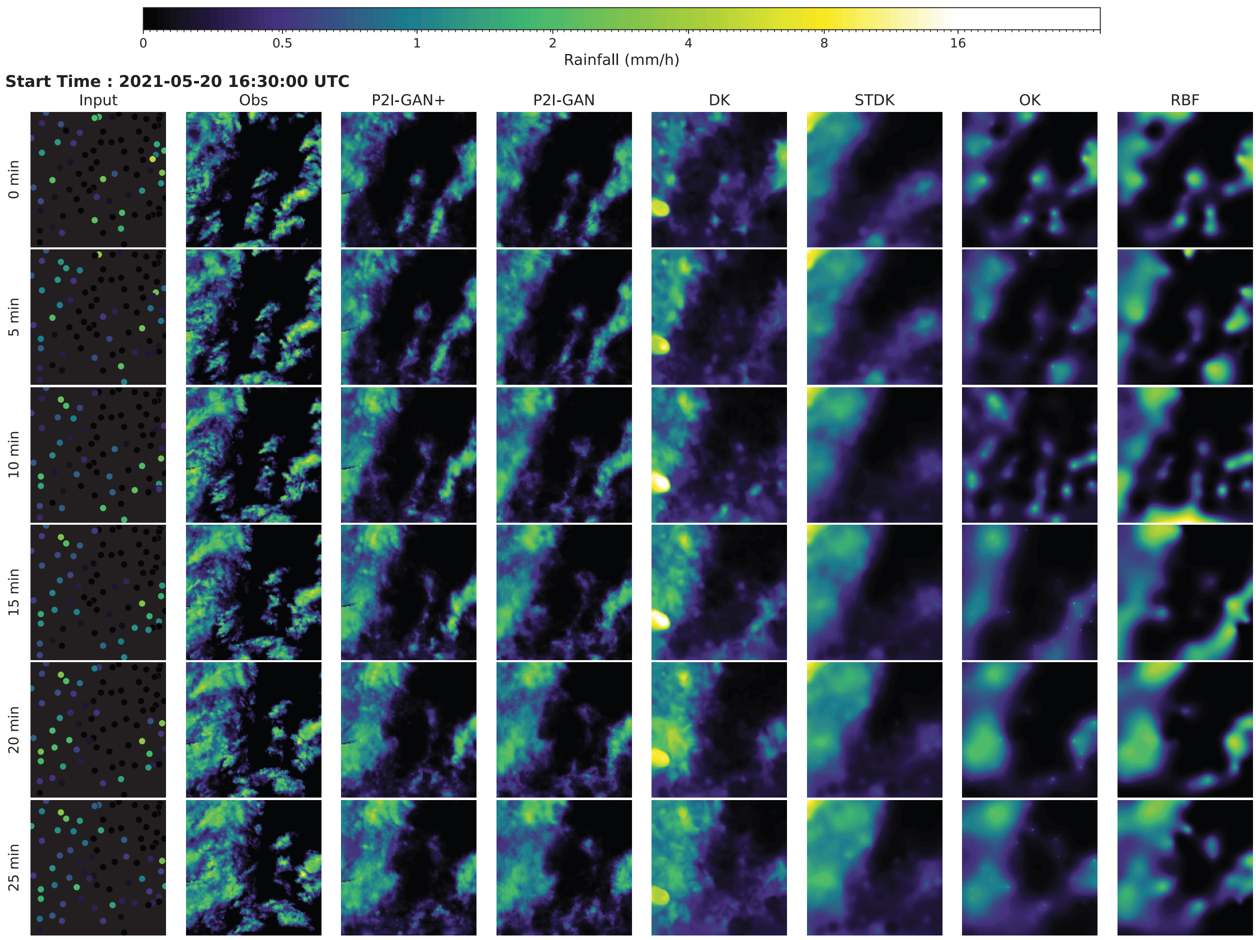

Figure S6. Representative interpolation case for Event 5 (2021-05-20 09:00 to 2021-05-22 02:00).

Event 6

Event 6 spans from 2021-05-23 11:00 to 2021-05-24 07:00, with a total duration of 20 hours. According to Table S4, the reported maximum rainfall intensities is 43.6 mm/h for the radar product (RD), while the corresponding mean intensities is 0.288 mm/h. Figure S7 presents one representative interpolation case for this event.

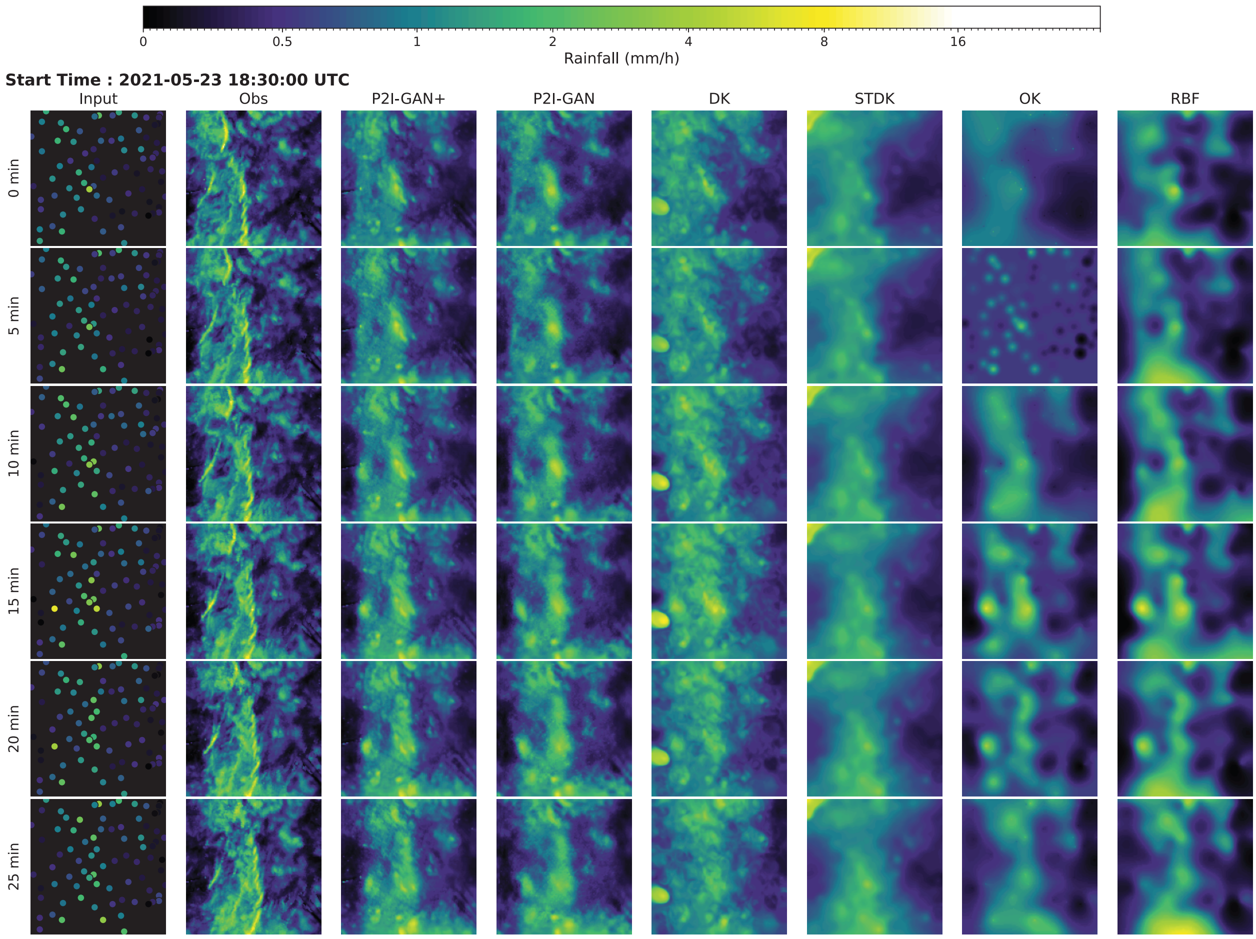

Figure S7. Representative interpolation case for Event 6 (2021-05-23 11:00 to 2021-05-24 07:00).

Event 7

Event 7 spans from 2021-07-05 17:00 to 2021-07-06 13:00, with a total duration of 20 hours. According to Table S4, the reported maximum rainfall intensities is 85.4 mm/h for the radar product (RD), while the corresponding mean intensities is 0.432 mm/h. Figure S8 presents one representative interpolation case for this event.

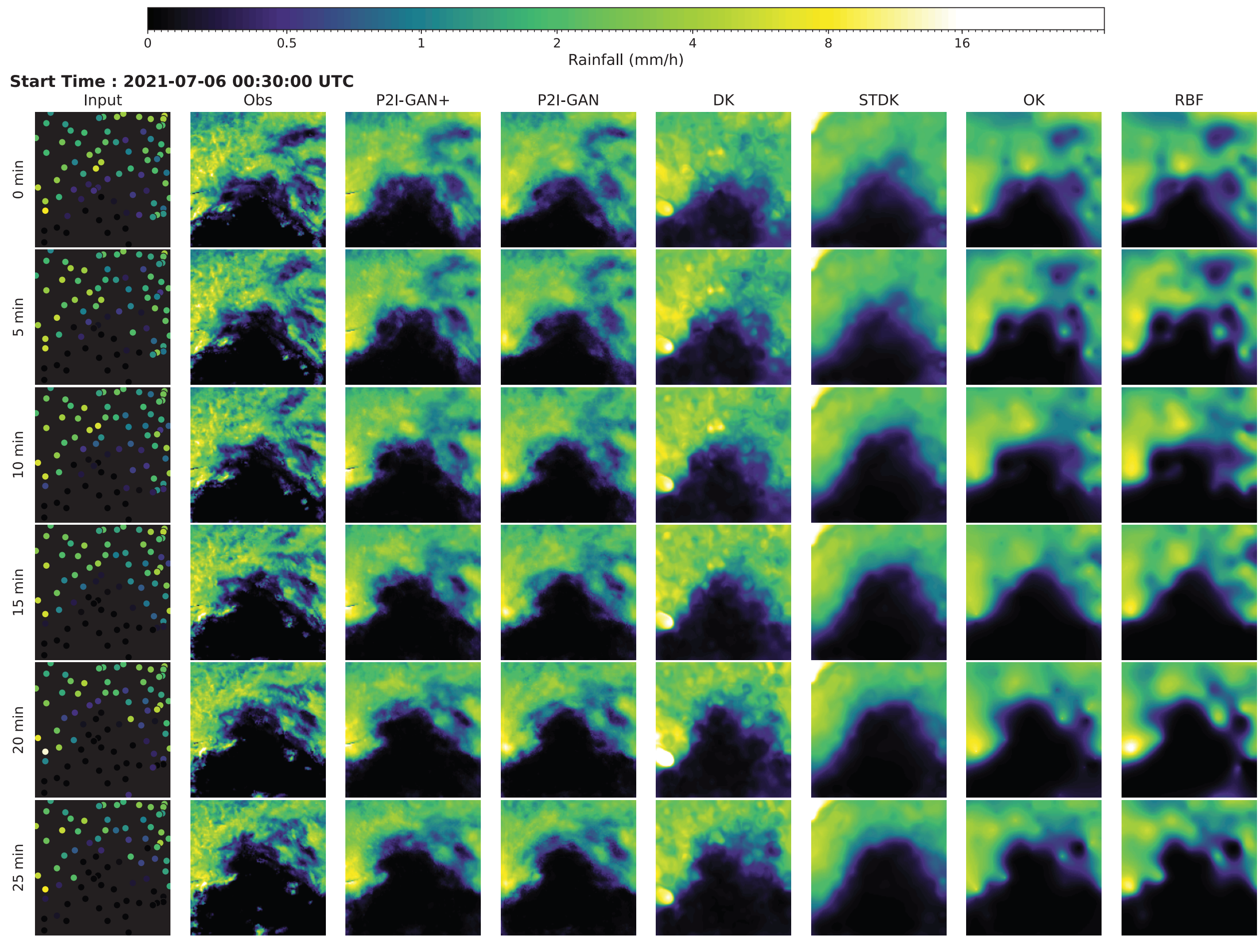

Figure S8. Representative interpolation case for Event 7 (2021-07-05 17:00 to 2021-07-06 13:00).

Event 8

Event 8 spans from 2021-10-02 05:00 to 2021-10-03 02:00, with a total duration of 21 hours. According to Table S4, the reported maximum rainfall intensities is 23.4 mm/h for the radar product (RD), while the corresponding mean intensities is 0.334 mm/h. Figure S9 presents one representative interpolation case for this event.

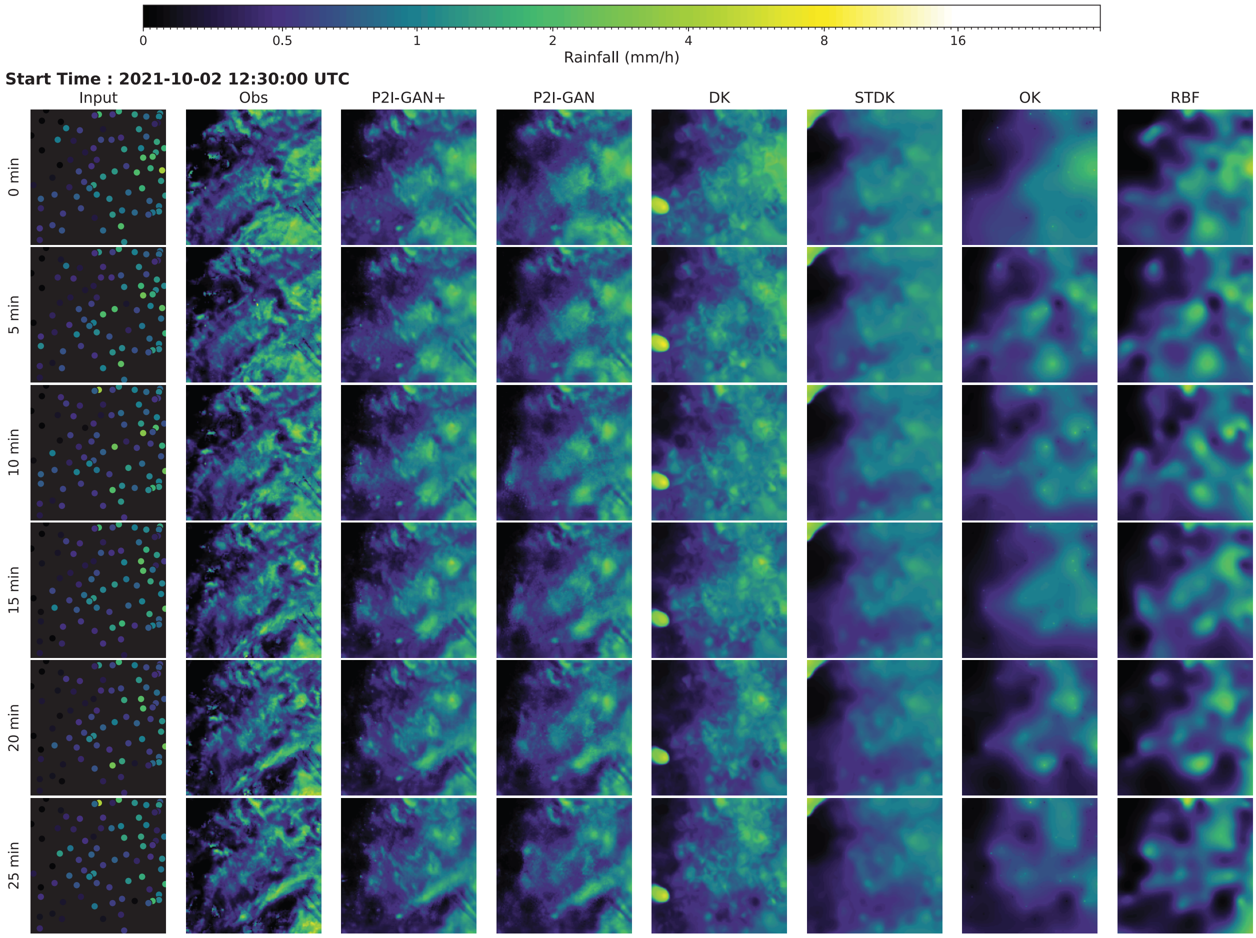

Figure S9. Representative interpolation case for Event 8 (2021-10-02 05:00 to 2021-10-03 02:00).

Event 9

Event 9 spans from 2021-10-31 03:00 to 2021-11-01 07:00, with a total duration of 28 hours. According to Table S4, the reported maximum rainfall intensities is 94.0 mm/h for the radar product (RD), while the corresponding mean intensities is 0.526 mm/h. Figure S10 presents one representative interpolation case for this event.

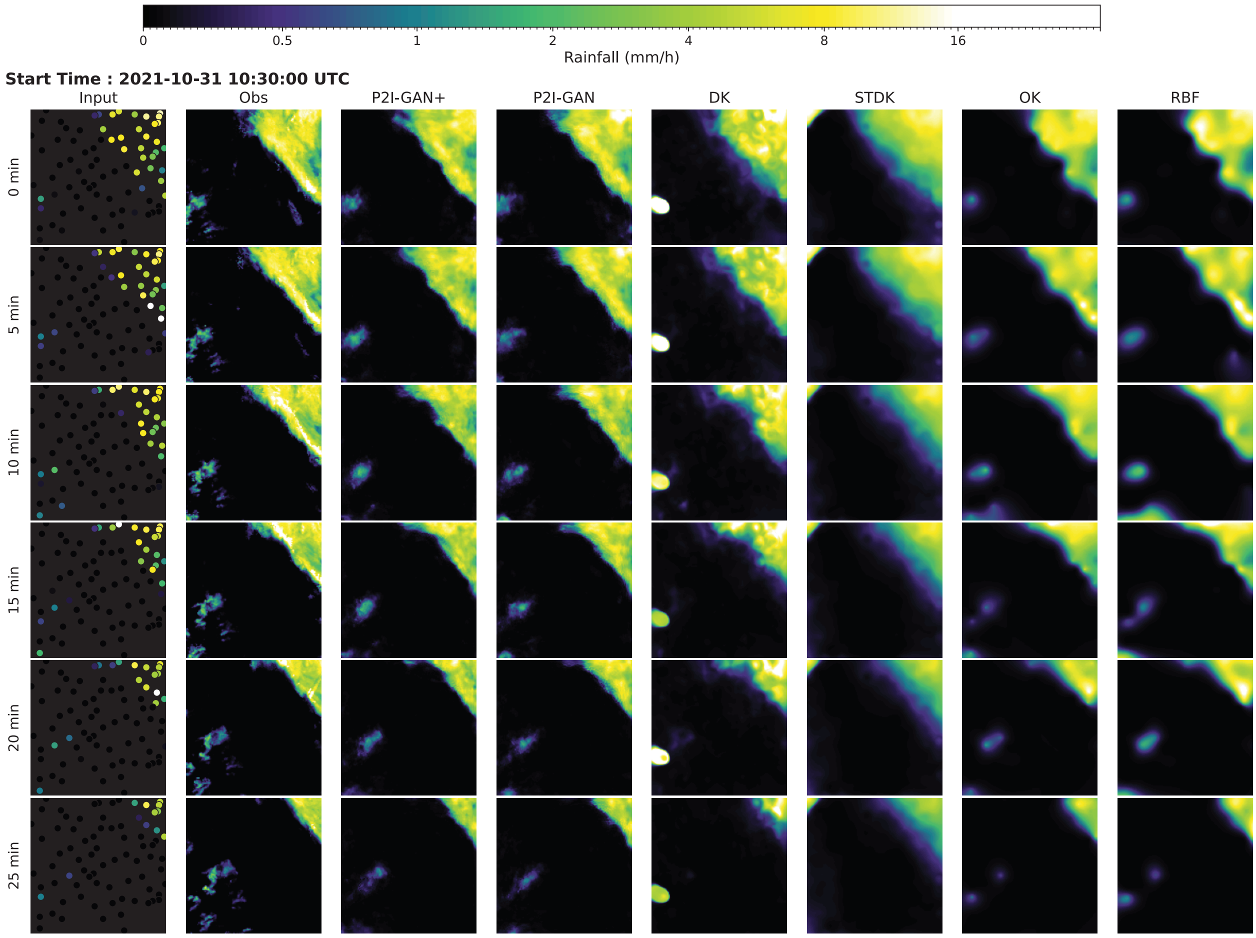

Figure S10. Representative interpolation case for Event 9 (2021-10-31 03:00 to 2021-11-01 07:00).

Event 10

Event 10 spans from 2021-12-07 08:00 to 2021-12-08 05:00, with a total duration of 21 hours. According to Table S4, the reported maximum rainfall intensities is 29.7 mm/h for the radar product (RD), while the corresponding mean intensities is 0.355 mm/h. Figure S11 presents one representative interpolation case for this event.

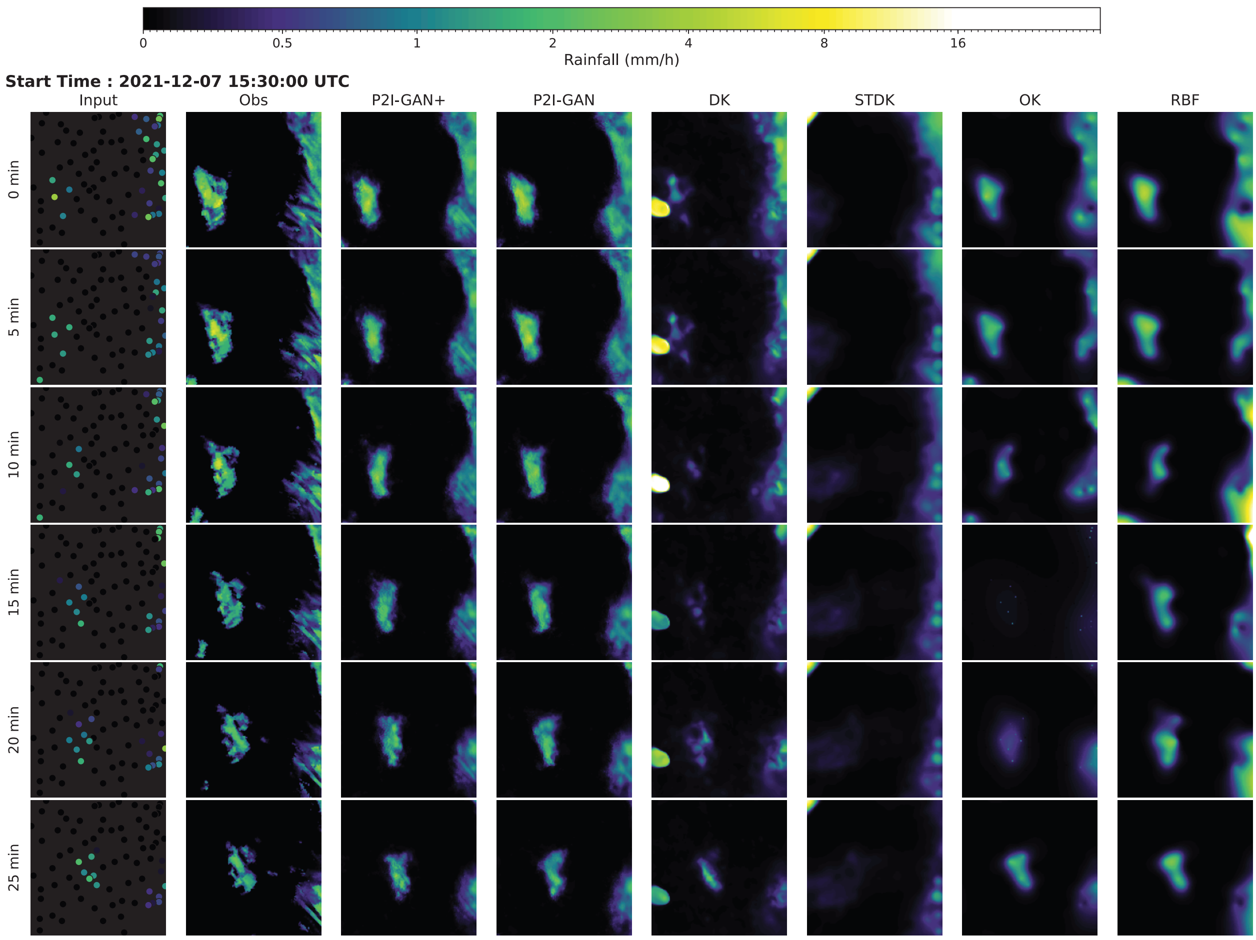

Figure S11. Representative interpolation case for Event 10 (2021-12-07 08:00 to 2021-12-08 05:00).

Event 11

Event 11 spans from 2022-01-08 04:00 to 2022-01-08 22:00, with a total duration of 18 hours. According to Table S4, the reported maximum rainfall intensities is 31.2 mm/h for the radar product (RD), while the corresponding mean intensities is 0.221 mm/h. Figure S12 presents one representative interpolation case for this event.

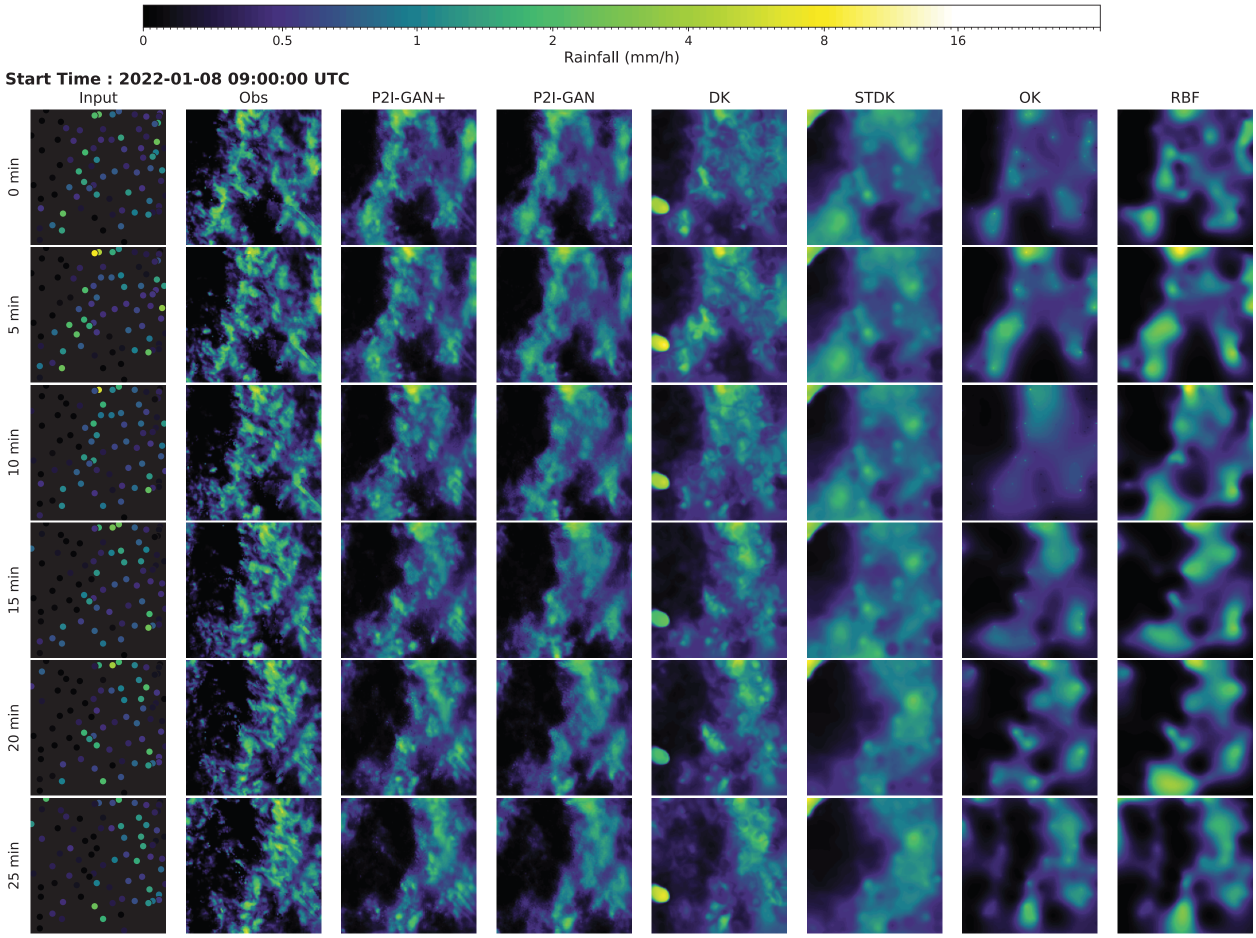

Figure S12. Representative interpolation case for Event 11 (2022-01-08 04:00 to 2022-01-08 22:00).

Event 12

Event 12 spans from 2022-03-16 11:00 to 2022-03-17 05:00, with a total duration of 18 hours. According to Table S4, the reported maximum rainfall intensities is 16.7 mm/h for the radar product (RD), while the corresponding mean intensities is 0.603 mm/h. Figure S13 presents one representative interpolation case for this event.

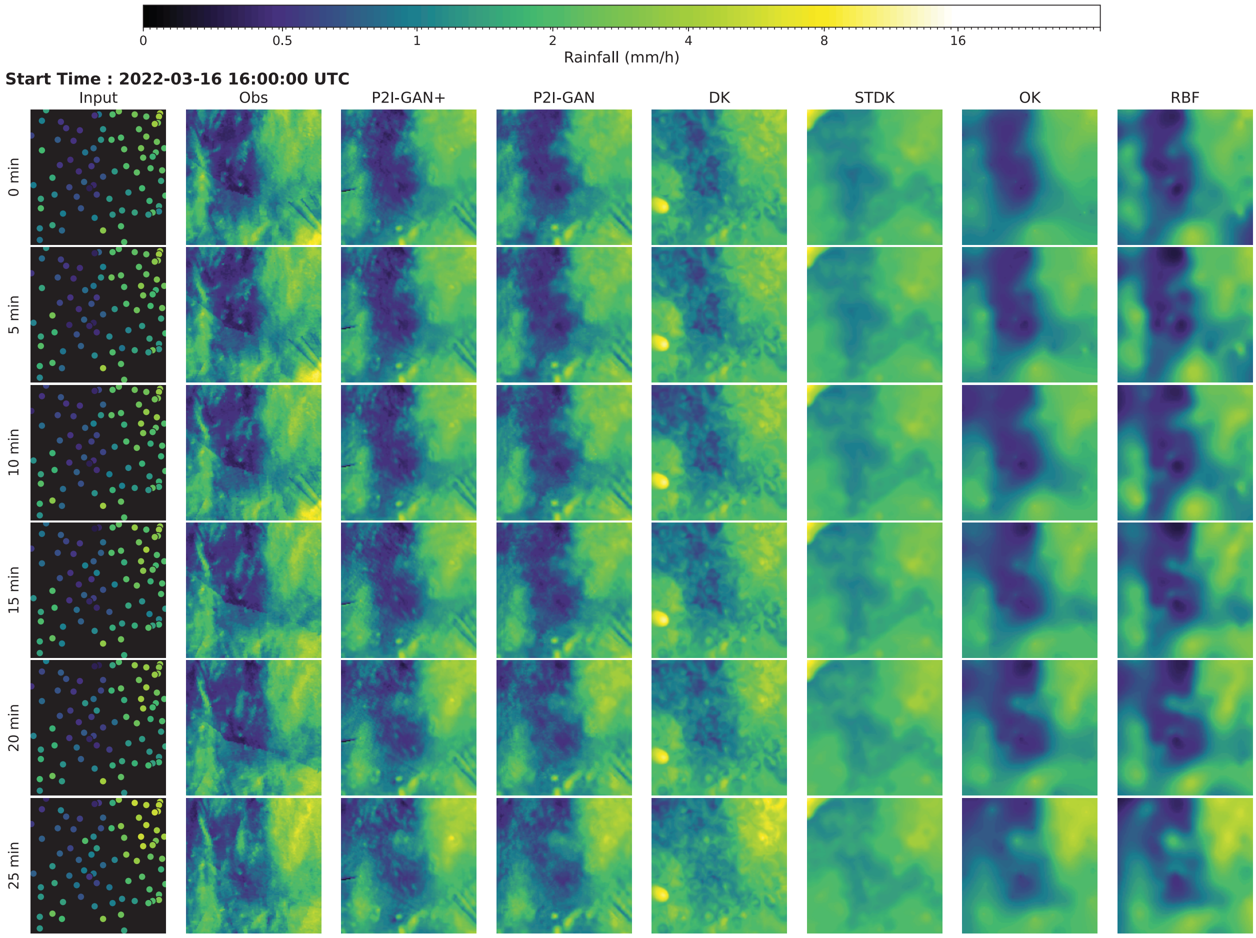

Figure S13. Representative interpolation case for Event 12 (2022-03-16 11:00 to 2022-03-17 05:00).

Event 13

Event 13 spans from 2022-04-12 06:00 to 2022-04-12 19:00, with a total duration of 13 hours. According to Table S4, the reported maximum rainfall intensities is 119.4 mm/h for the radar product (RD), while the corresponding mean intensities is 0.337 mm/h. Figure S14 presents one representative interpolation case for this event.

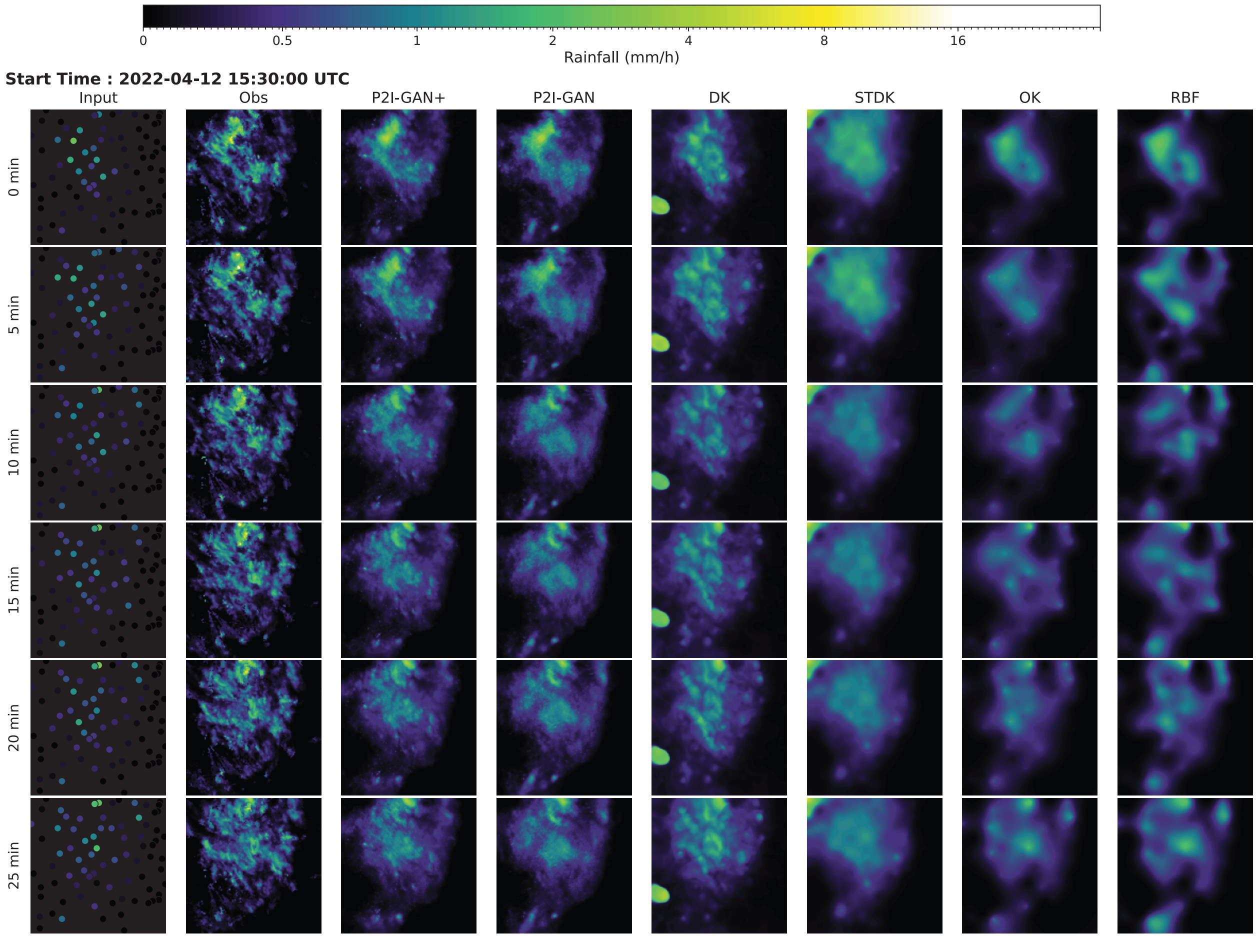

Figure S14. Representative interpolation case for Event 13 (2022-04-12 06:00 to 2022-04-12 19:00).

Event 14

Event 14 spans from 2022-05-11 02:00 to 2022-05-11 20:00, with a total duration of 18 hours. According to Table S4, the reported maximum rainfall intensities is 85.4 mm/h for the radar product (RD), while the corresponding mean intensities is 0.417 mm/h. Figure S15 presents one representative interpolation case for this event.

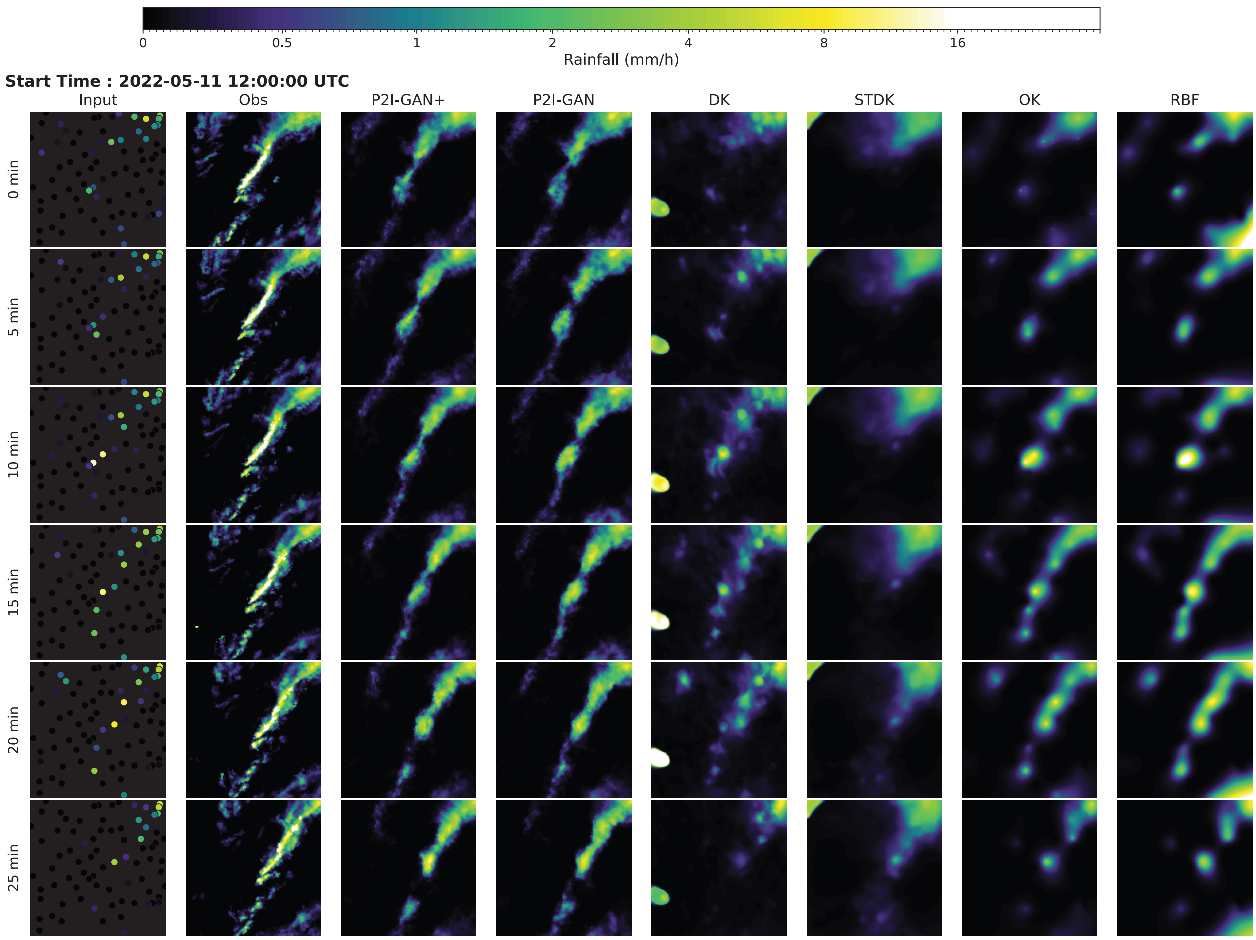

Figure S15. Representative interpolation case for Event 14 (2022-05-11 02:00 to 2022-05-11 20:00).

Event 15

Event 15 spans from 2022-05-15 20:00 to 2022-05-16 11:00, with a total duration of 15 hours. According to Table S4, the reported maximum rainfall intensities is 89.6 mm/h for the radar product (RD), while the corresponding mean intensities is 0.287 mm/h. Figure S16 presents one representative interpolation case for this event.

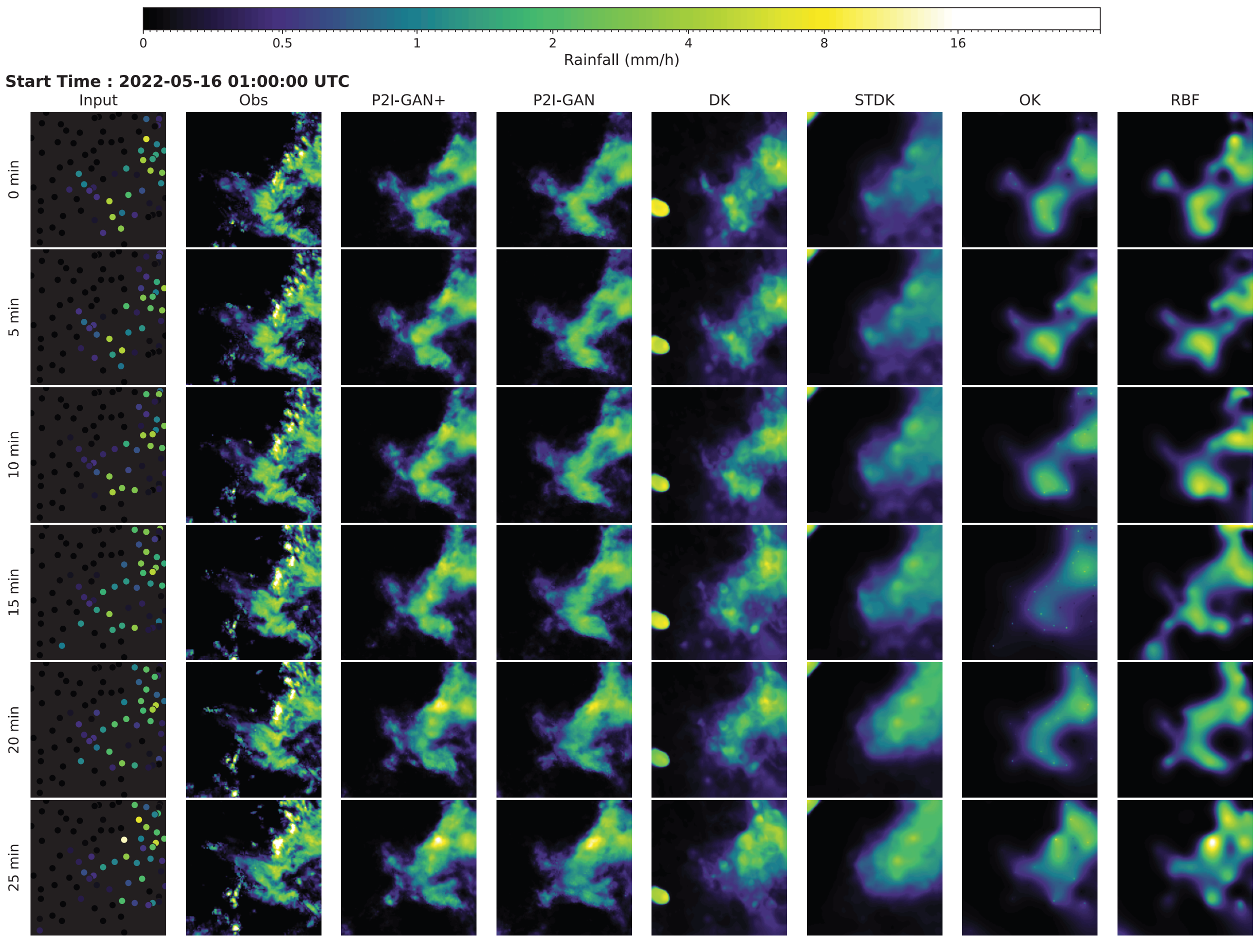

Figure S16. Representative interpolation case for Event 15 (2022-05-15 20:00 to 2022-05-16 11:00).

Event 16

Event 16 spans from 2022-06-18 15:00 to 2022-06-19 06:00, with a total duration of 15 hours. According to Table S4, the reported maximum rainfall intensities is 55.4 mm/h for the radar product (RD), while the corresponding mean intensities is 0.459 mm/h. Figure S17 presents one representative interpolation case for this event.

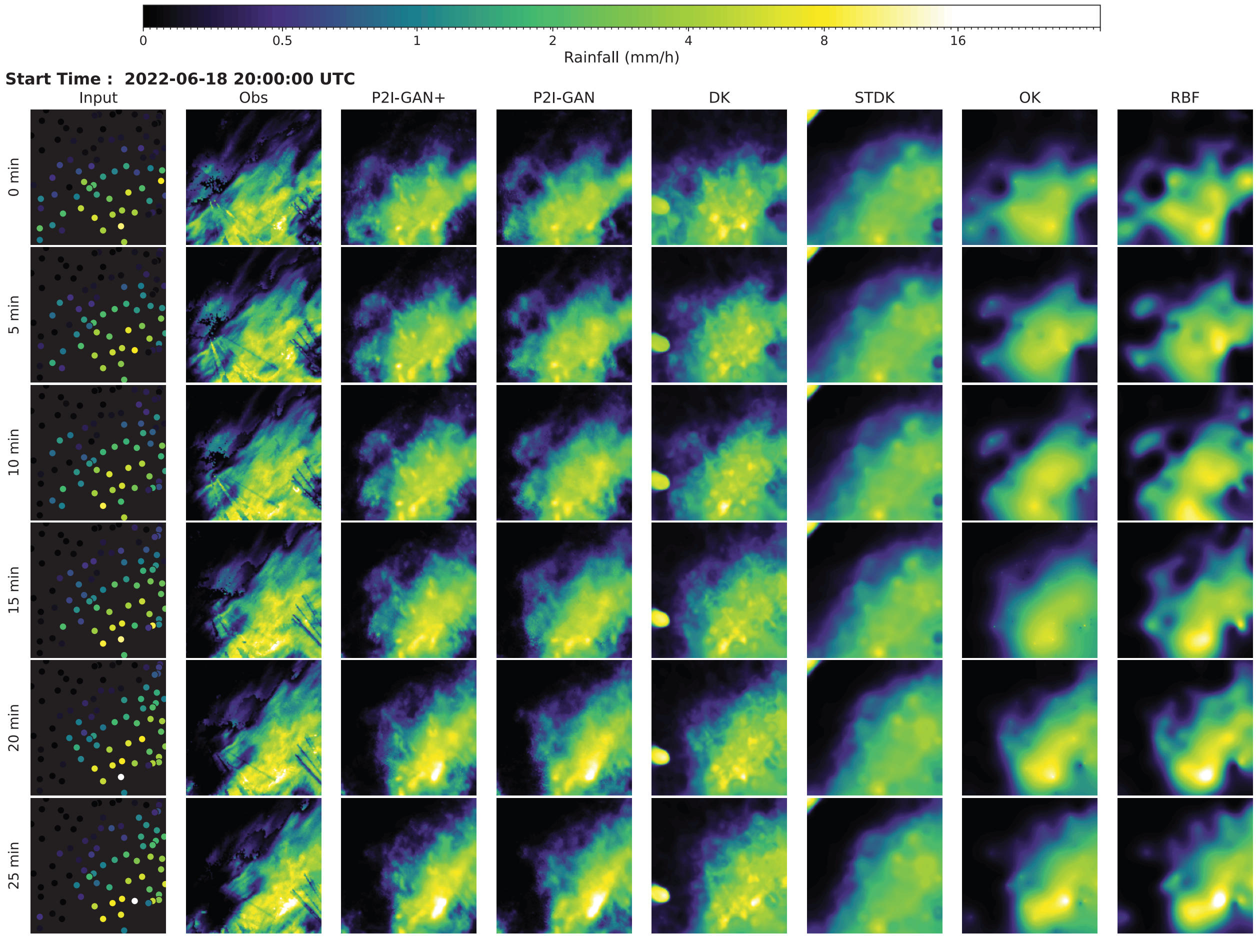

Figure S17. Representative interpolation case for Event 16 (2022-06-18 15:00 to 2022-06-19 06:00).

Event 17

Event 17 spans from 2022-08-25 05:00 to 2022-08-25 19:00, with a total duration of 14 hours. According to Table S4, the reported maximum rainfall intensities is 45.8 mm/h for the radar product (RD), while the corresponding mean intensities is 0.232 mm/h. Figure S18 presents one representative interpolation case for this event.

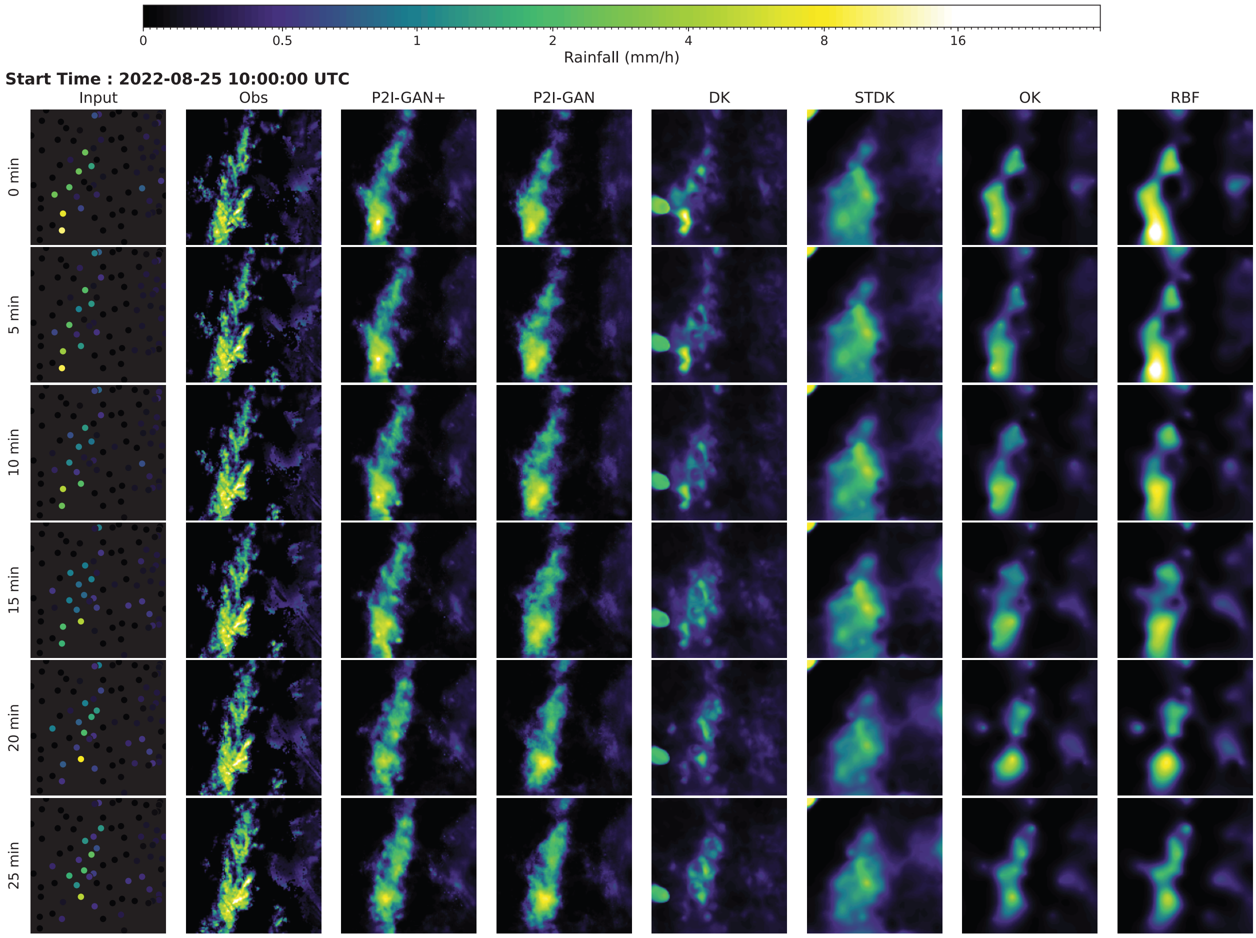

Figure S18. Representative interpolation case for Event 17 (2022-08-25 05:00 to 2022-08-25 19:00).

Event 18

Event 18 spans from 2022-09-30 09:00 to 2022-10-01 03:00, with a total duration of 18 hours. According to Table S4, the reported maximum rainfall intensities is 31.2 mm/h for the radar product (RD), while the corresponding mean intensities is 0.474 mm/h. Figure S19 presents one representative interpolation case for this event.

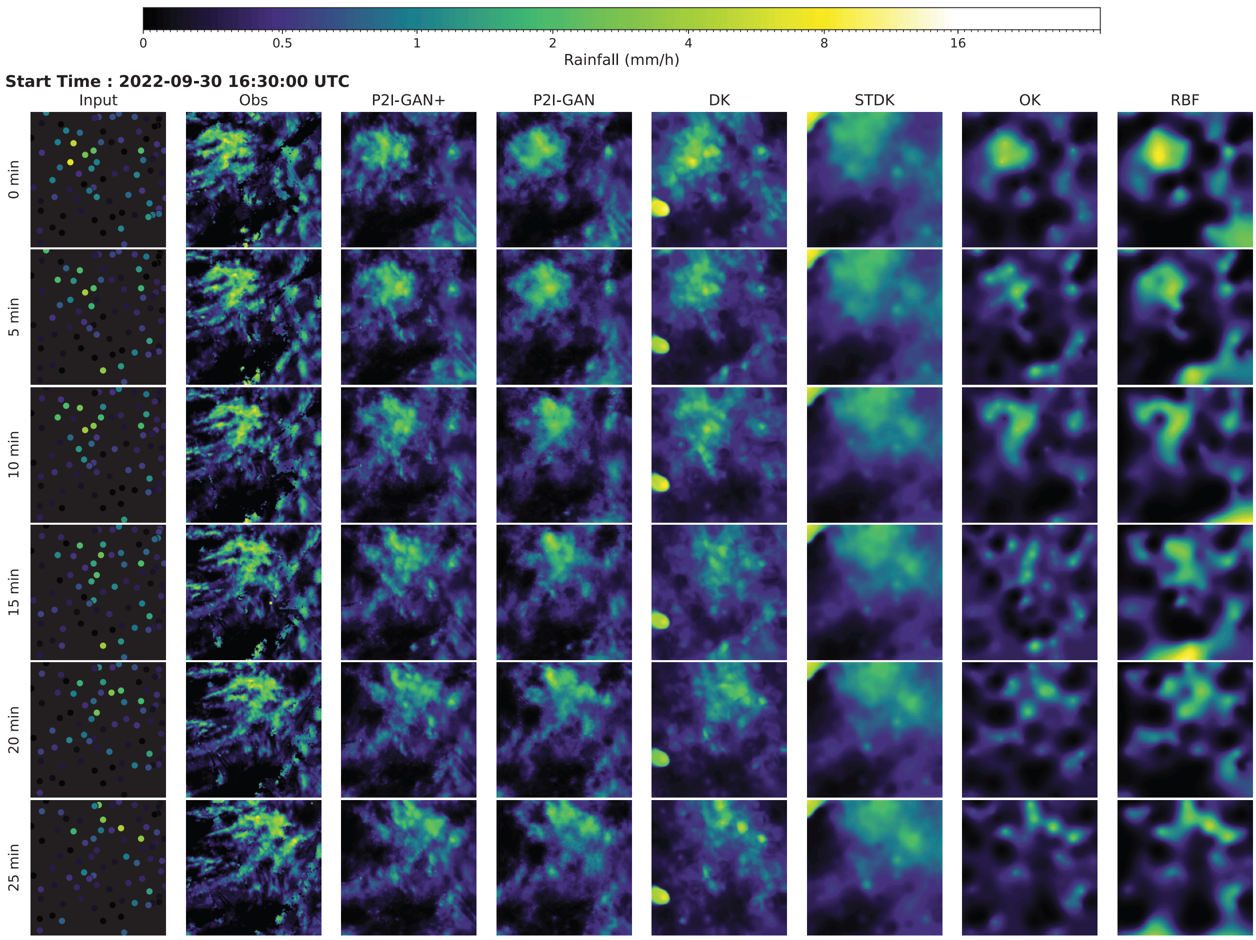

Figure S19. Representative interpolation case for Event 18 (2022-09-30 09:00 to 2022-10-01 03:00).

Event 19

Event 19 spans from 2022-10-31 18:00 to 2022-11-01 12:00, with a total duration of 18 hours. According to Table S4, the reported maximum rainfall intensities is 77.6 mm/h for the radar product (RD), while the corresponding mean intensities is 0.580 mm/h. Figure S20 presents one representative interpolation case for this event.

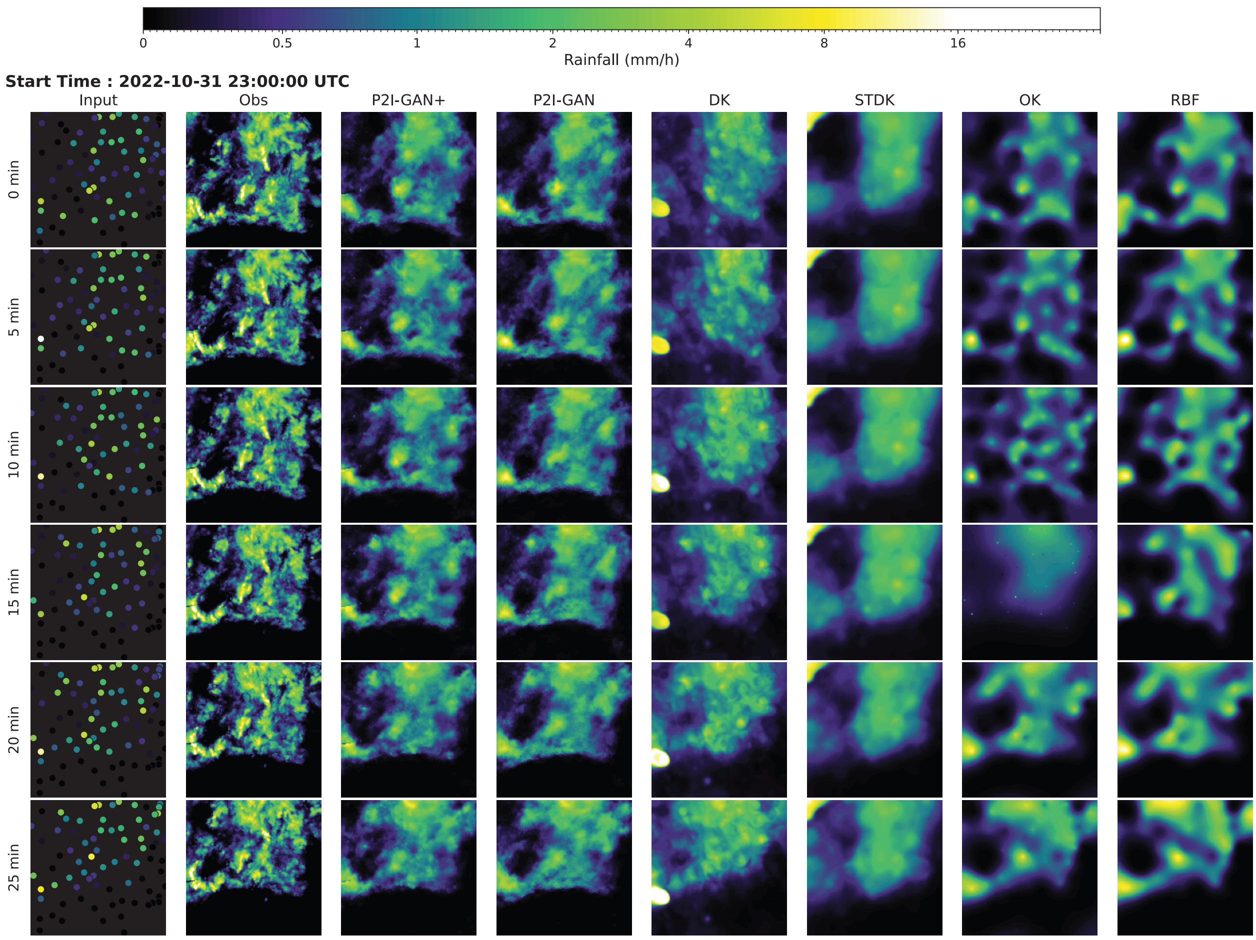

Figure S20. Representative interpolation case for Event 19 (2022-10-31 18:00 to 2022-11-01 12:00).

Event 20

Event 20 spans from 2022-11-23 03:00 to 2022-11-23 19:00, with a total duration of 16 hours. According to Table S4, the reported maximum rainfall intensities is 85.4 mm/h for the radar product (RD), while the corresponding mean intensities is 0.506 mm/h. Figure S21 presents one representative interpolation case for this event.

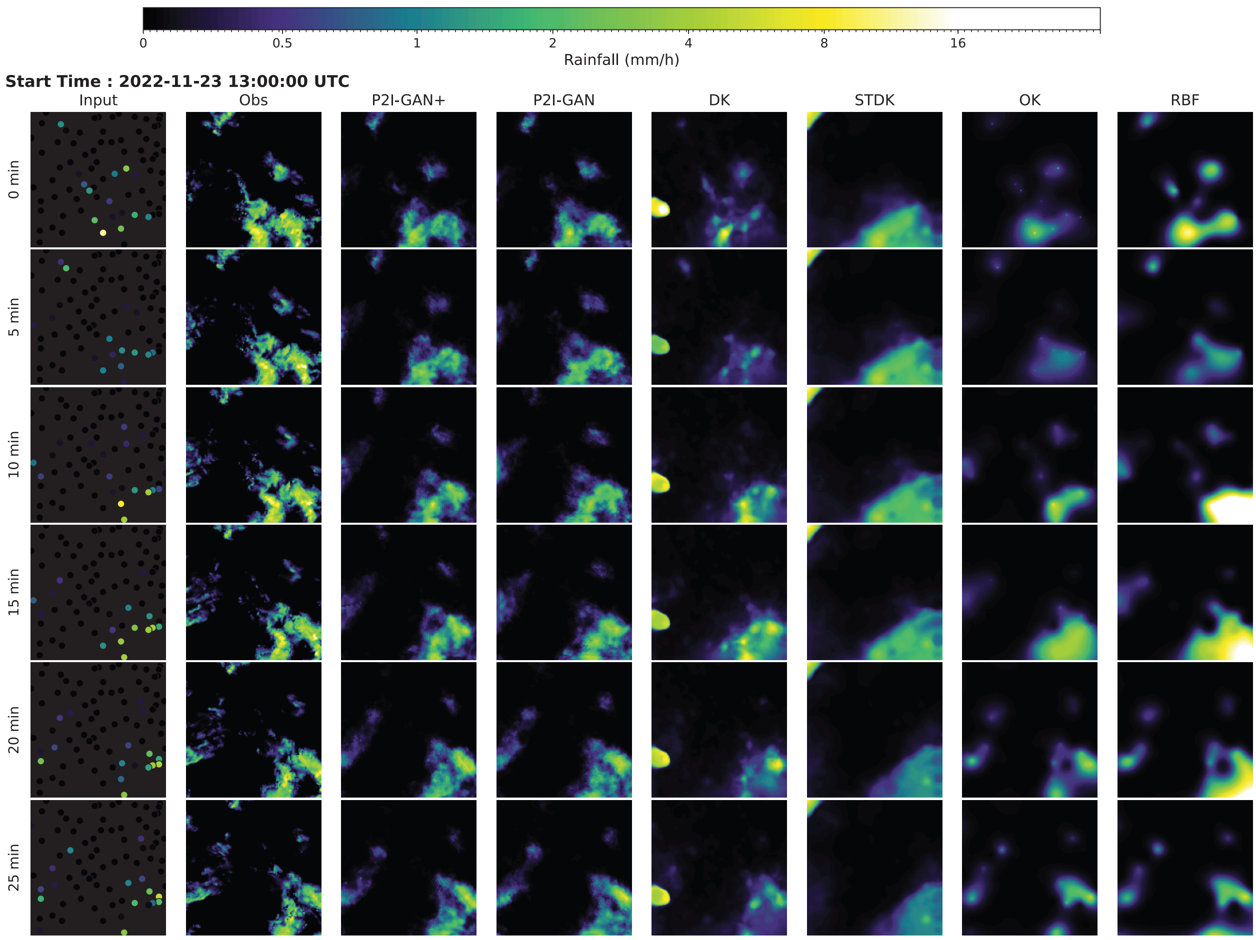

Figure S21. Representative interpolation case for Event 20 (2022-11-23 03:00 to 2022-11-23 19:00).

S2.2 Stage II: 10 event cases

The following cases provide representative gauge-derived reconstruction figures for the first ten selected events.

Event 1

Event 1 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-01-15 22:00 to 2021-01-16 18:00 and lasts 20 hours, with reported maximum rainfall intensities of 16.8 mm/h (RG) and 48.0 mm/h (RD), and reported mean intensities of 0.313 mm/h (RG) and 0.424 mm/h (RD). Figure S22 provides the corresponding representative reconstruction example.

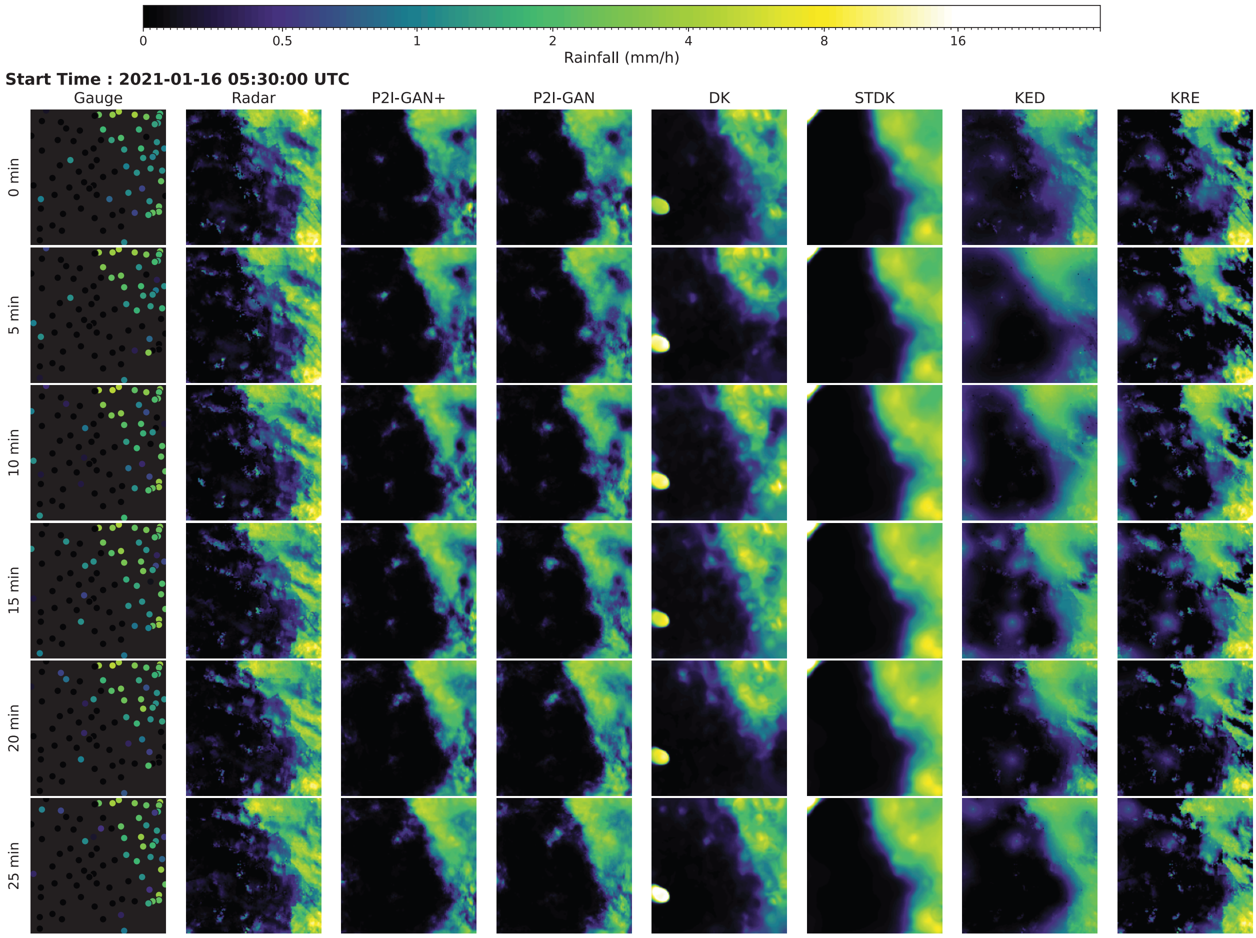

Figure S22. Representative gauge-derived reconstruction case for Event 1 (2021-01-15 22:00 to 2021-01-16 18:00).

Event 2

Event 2 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-03-09 22:00 to 2021-03-11 00:00 and lasts 26 hours, with reported maximum rainfall intensities of 17.4 mm/h (RG) and 13.8 mm/h (RD), and reported mean intensities of 0.220 mm/h (RG) and 0.285 mm/h (RD). Figure S23 provides the corresponding representative reconstruction example.

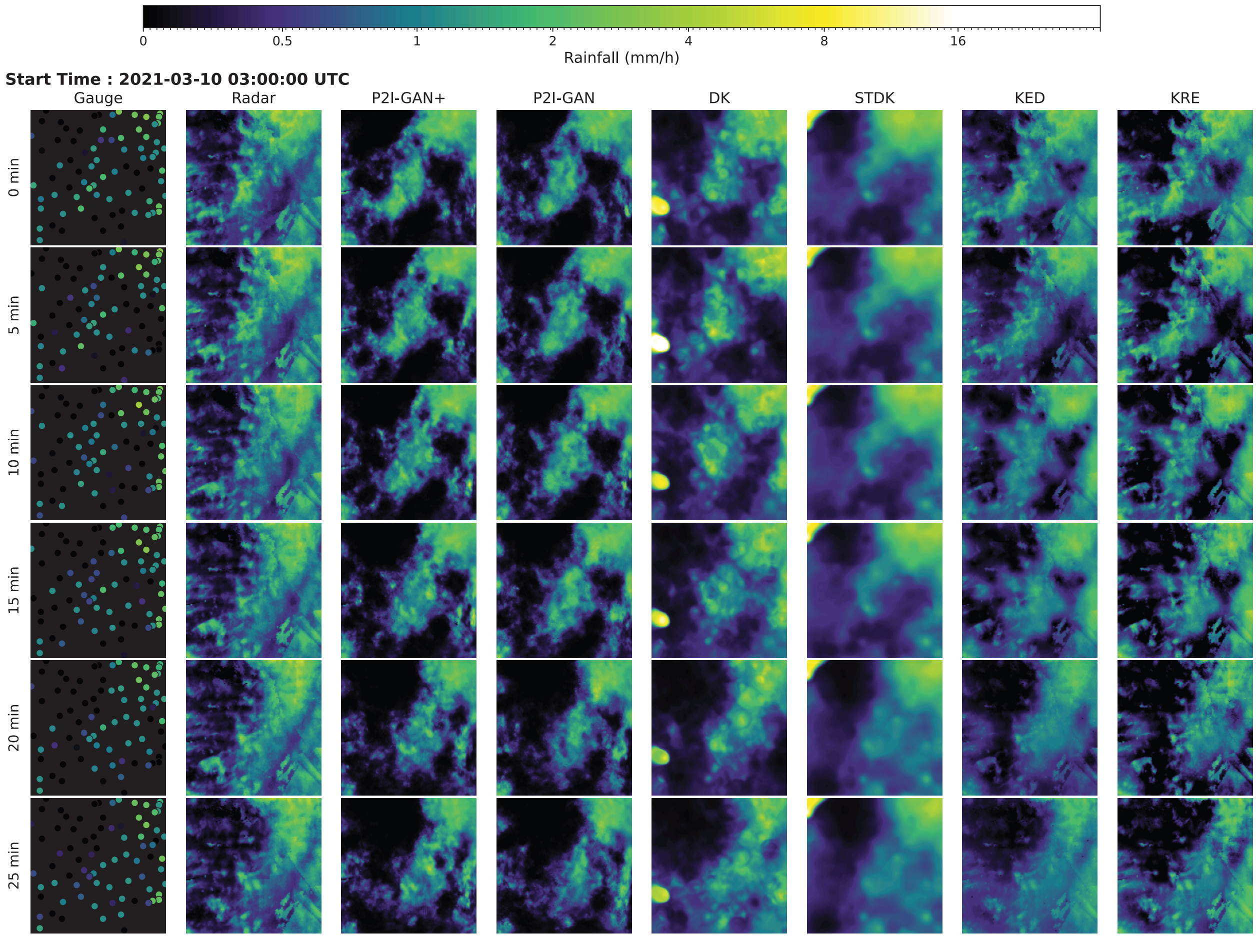

Figure S23. Representative gauge-derived reconstruction case for Event 2 (2021-03-09 22:00 to 2021-03-11 00:00).

Event 3

Event 3 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-05-03 08:00 to 2021-05-04 05:00 and lasts 21 hours, with reported maximum rainfall intensities of 47.3 mm/h (RG) and 45.8 mm/h (RD), and reported mean intensities of 0.722 mm/h (RG) and 0.642 mm/h (RD). Figure S24 provides the corresponding representative reconstruction example.

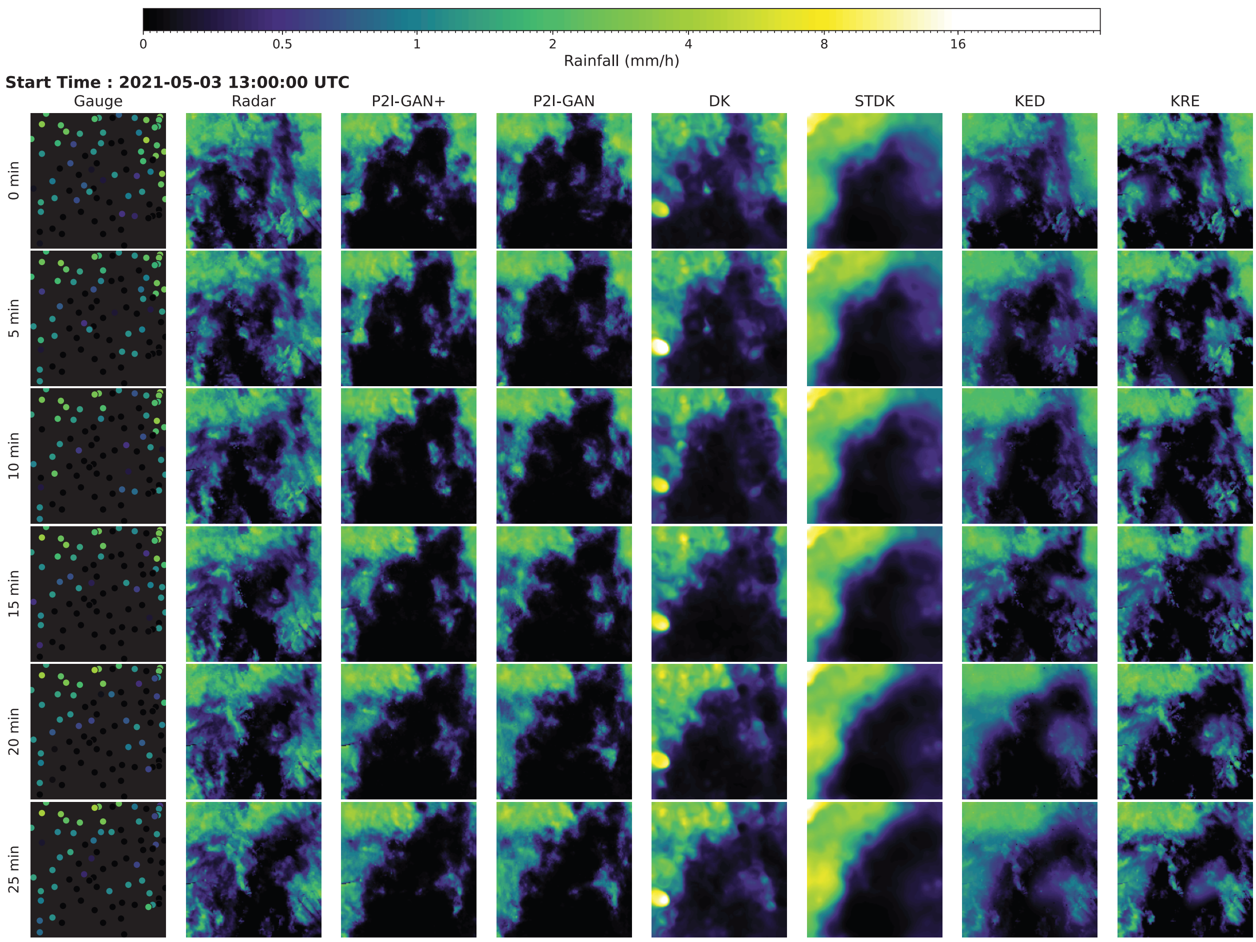

Figure S24. Representative gauge-derived reconstruction case for Event 3 (2021-05-03 08:00 to 2021-05-04 05:00).

Event 4

Event 4 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-05-08 01:00 to 2021-05-08 21:00 and lasts 20 hours, with reported maximum rainfall intensities of 13.4 mm/h (RG) and 21.2 mm/h (RD), and reported mean intensities of 0.878 mm/h (RG) and 0.670 mm/h (RD). Figure S25 provides the corresponding representative reconstruction example.

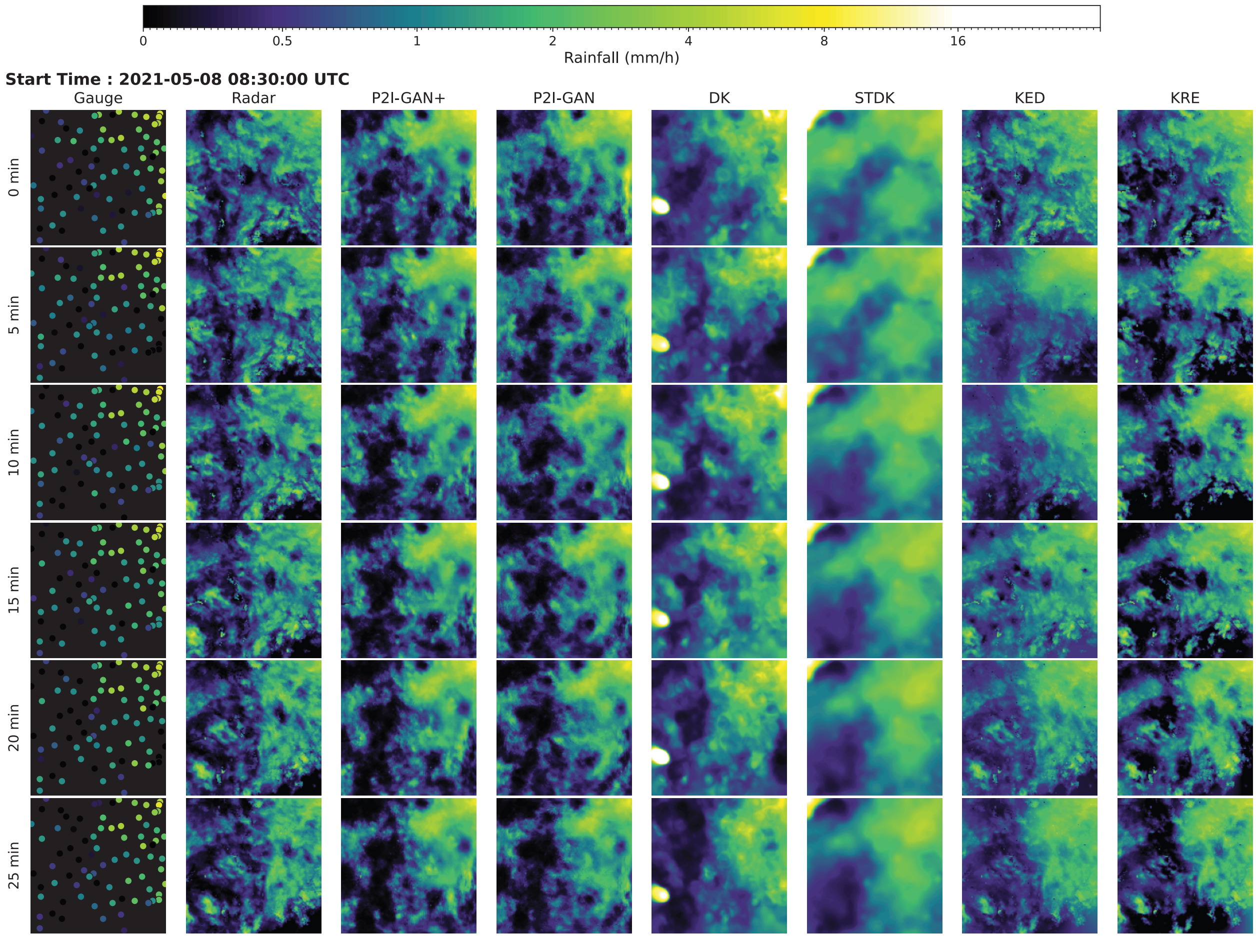

Figure S25. Representative gauge-derived reconstruction case for Event 4 (2021-05-08 01:00 to 2021-05-08 21:00).

Event 5

Event 5 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-05-20 09:00 to 2021-05-22 02:00 and lasts 41 hours, with reported maximum rainfall intensities of 21.0 mm/h (RG) and 32.7 mm/h (RD), and reported mean intensities of 0.368 mm/h (RG) and 0.386 mm/h (RD). Figure S26 provides the corresponding representative reconstruction example.

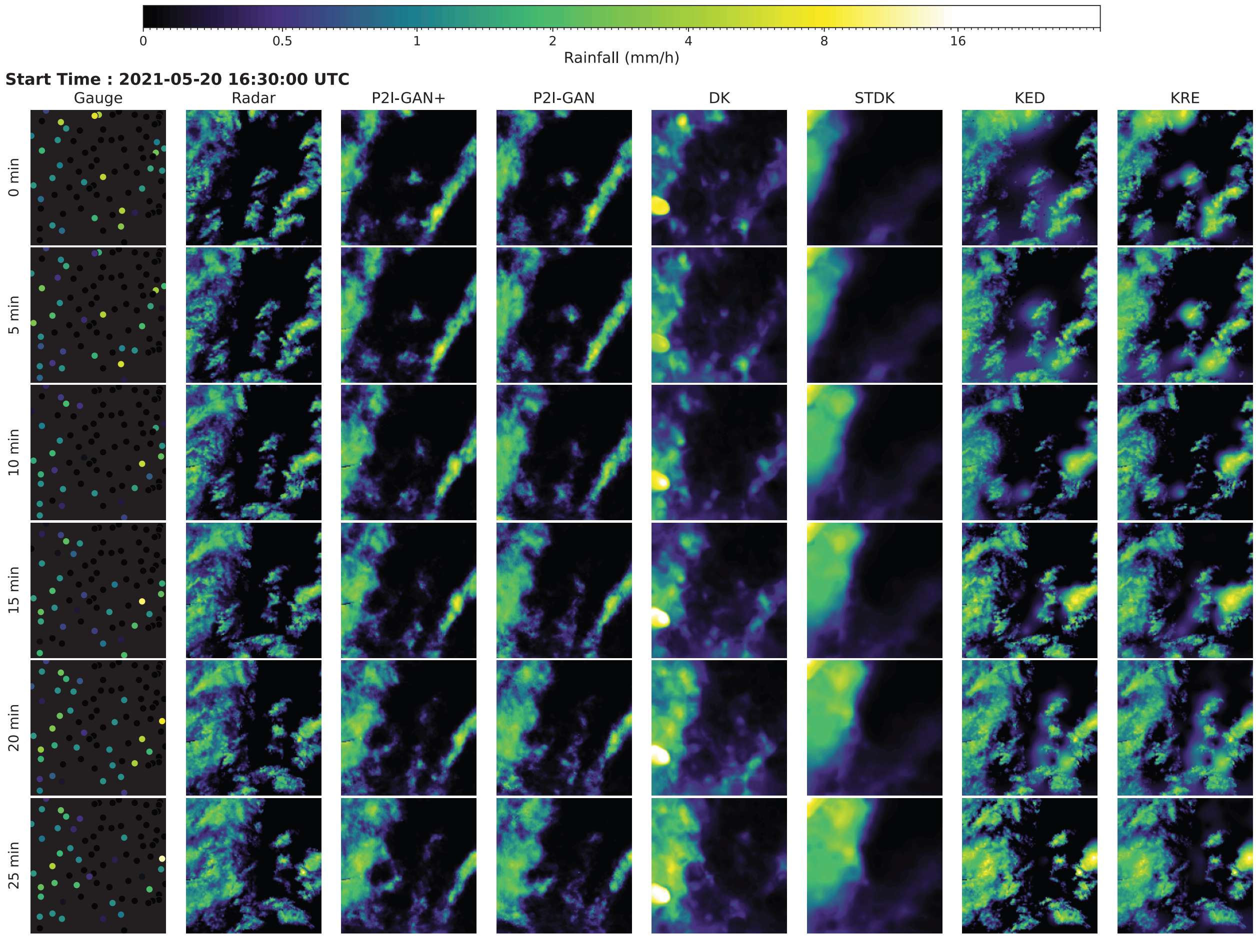

Figure S26. Representative gauge-derived reconstruction case for Event 5 (2021-05-20 09:00 to 2021-05-22 02:00).

Event 6

Event 6 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-05-23 11:00 to 2021-05-24 07:00 and lasts 20 hours, with reported maximum rainfall intensities of 34.8 mm/h (RG) and 43.6 mm/h (RD), and reported mean intensities of 0.325 mm/h (RG) and 0.288 mm/h (RD). Figure S27 provides the corresponding representative reconstruction example.

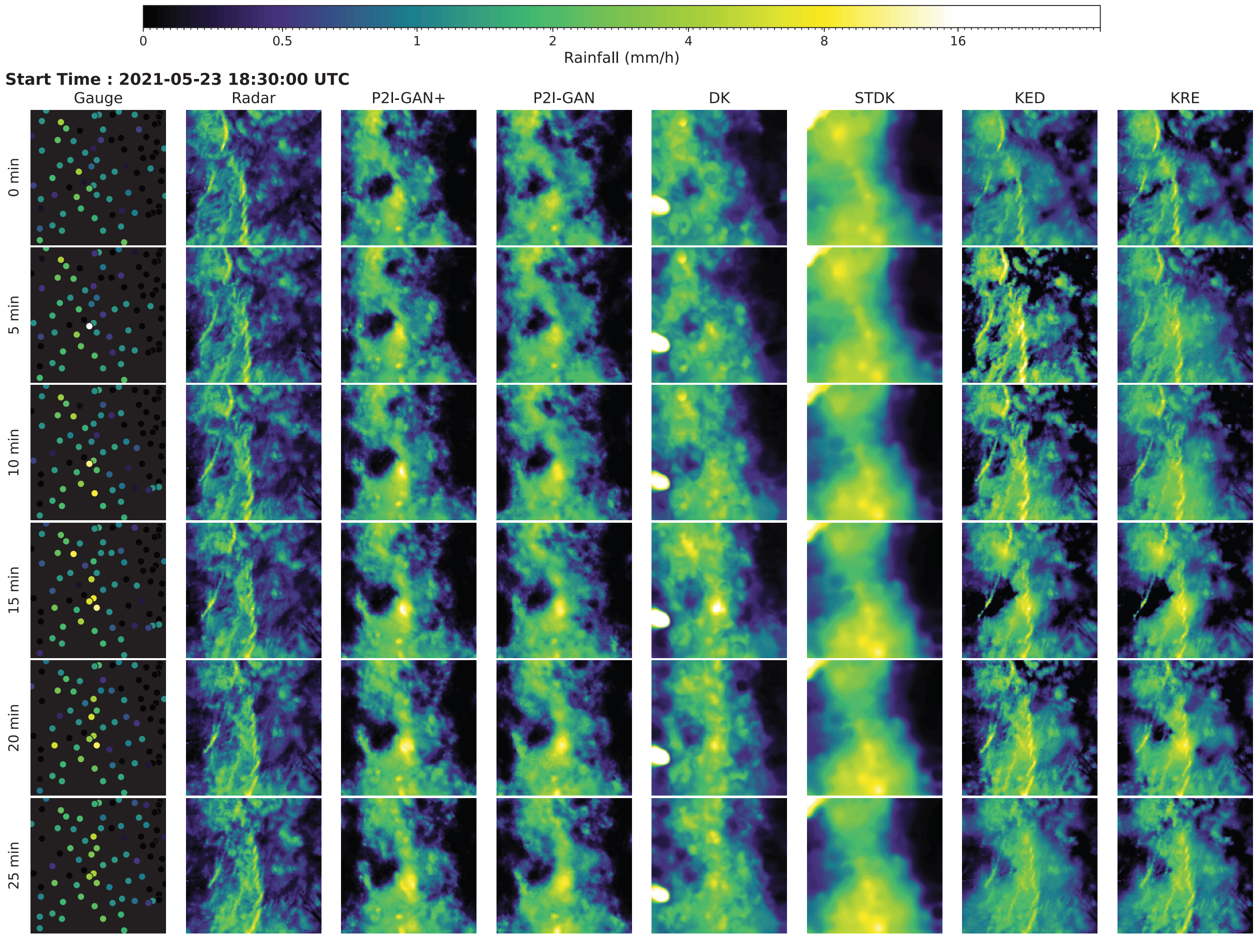

Figure S27. Representative gauge-derived reconstruction case for Event 6 (2021-05-23 11:00 to 2021-05-24 07:00).

Event 7

Event 7 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-07-05 17:00 to 2021-07-06 13:00 and lasts 20 hours, with reported maximum rainfall intensities of 22.6 mm/h (RG) and 85.4 mm/h (RD), and reported mean intensities of 0.530 mm/h (RG) and 0.432 mm/h (RD). Figure S28 provides the corresponding representative reconstruction example.

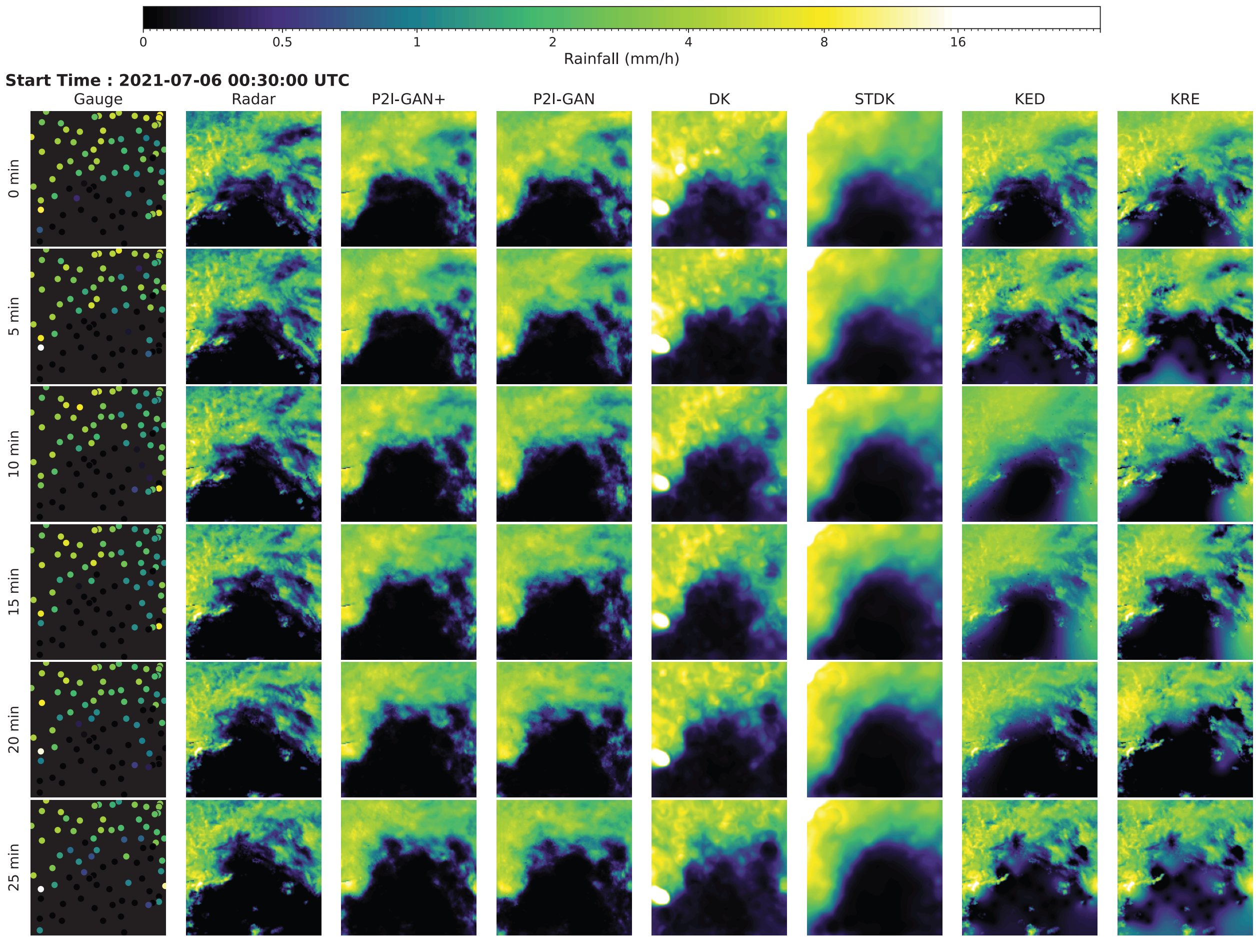

Figure S28. Representative gauge-derived reconstruction case for Event 7 (2021-07-05 17:00 to 2021-07-06 13:00).

Event 8

Event 8 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-10-02 05:00 to 2021-10-03 02:00 and lasts 21 hours, with reported maximum rainfall intensities of 13.0 mm/h (RG) and 23.4 mm/h (RD), and reported mean intensities of 0.377 mm/h (RG) and 0.334 mm/h (RD). Figure S29 provides the corresponding representative reconstruction example.

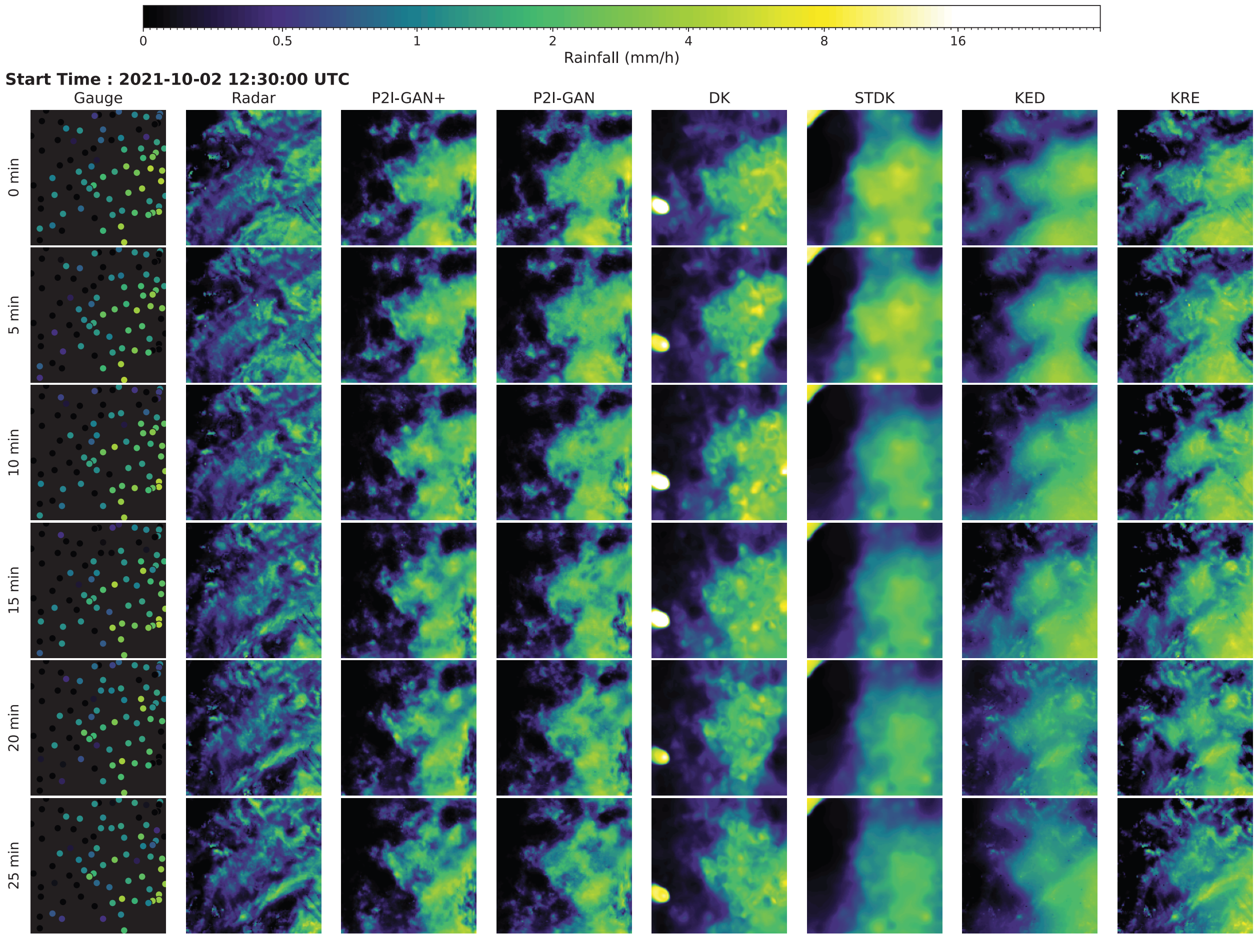

Figure S29. Representative gauge-derived reconstruction case for Event 8 (2021-10-02 05:00 to 2021-10-03 02:00).

Event 9

Event 9 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-10-31 03:00 to 2021-11-01 07:00 and lasts 28 hours, with reported maximum rainfall intensities of 89.5 mm/h (RG) and 94.0 mm/h (RD), and reported mean intensities of 0.627 mm/h (RG) and 0.526 mm/h (RD). Figure S30 provides the corresponding representative reconstruction example.

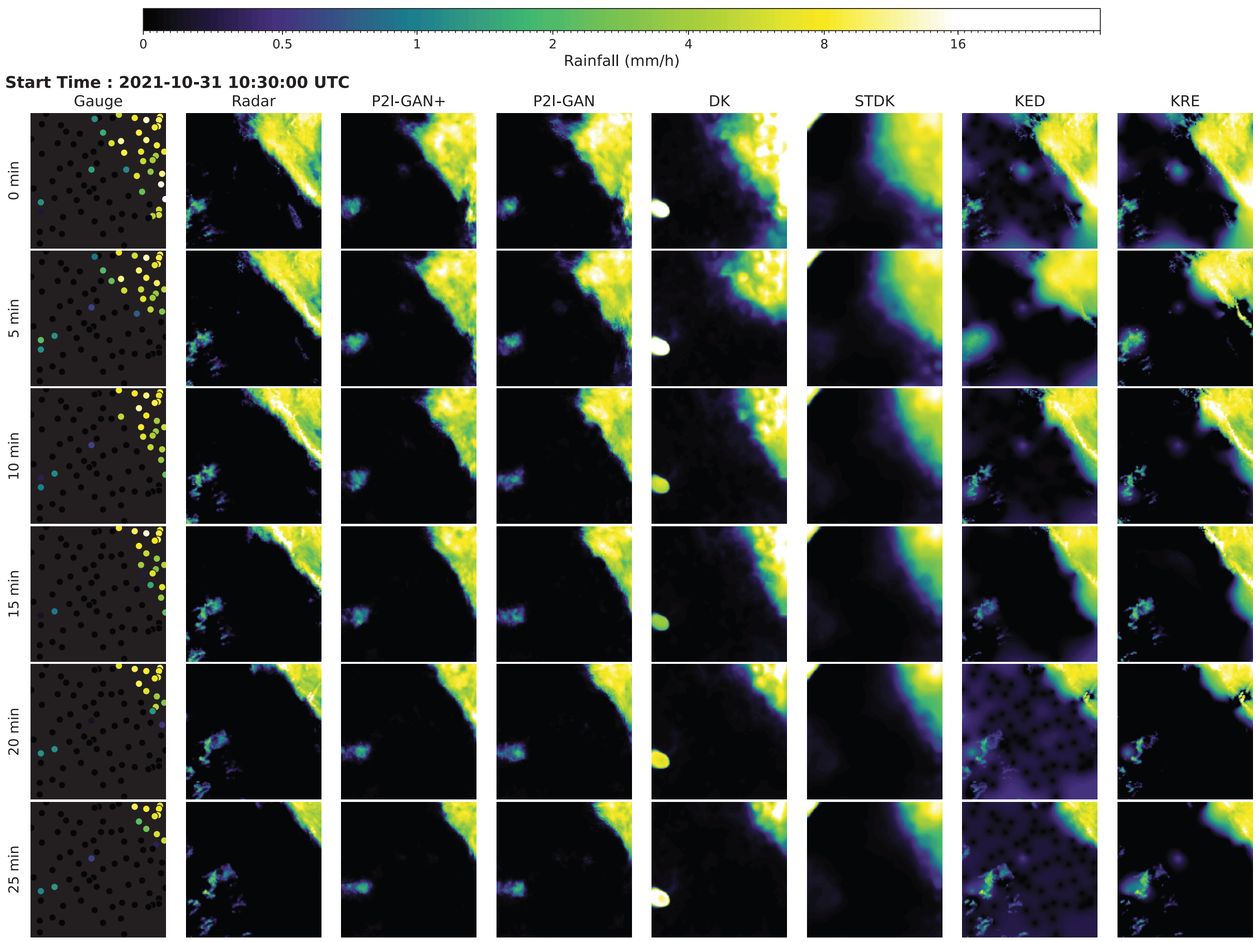

Figure S30. Representative gauge-derived reconstruction case for Event 9 (2021-10-31 03:00 to 2021-11-01 07:00).

Event 10

Event 10 is included in the first ten gauge-derived reconstruction cases. The event extends from 2021-12-07 08:00 to 2021-12-08 05:00 and lasts 21 hours, with reported maximum rainfall intensities of 39.2 mm/h (RG) and 29.7 mm/h (RD), and reported mean intensities of 0.293 mm/h (RG) and 0.355 mm/h (RD). Figure S31 provides the corresponding representative reconstruction example.

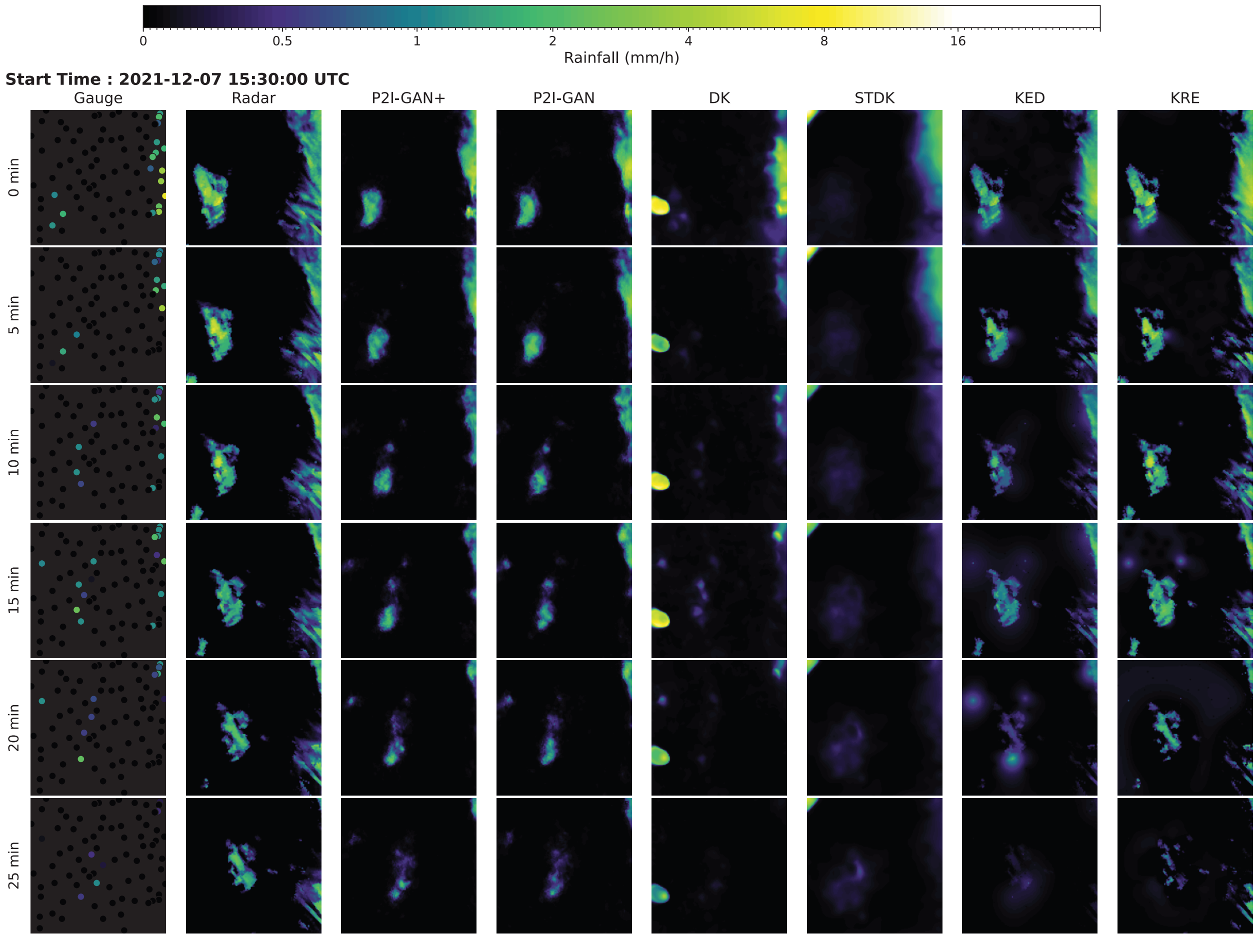

Figure S31. Representative gauge-derived reconstruction case for Event 10 (2021-12-07 08:00 to 2021-12-08 05:00).

References

Chen, W., Li, Y., Reich, B. J., and Sun, Y.: DeepKriging: Spatially Dependent Deep Neural Networks for Spatial Prediction, https://arxiv.org/abs/2007.11972, 2022.

Nag, P., Sun, Y., and Reich, B. J.: Spatio-temporal DeepKriging for interpolation and probabilistic forecasting, Spatial Statistics, 57, 100 773, https://doi.org/https://doi.org/10.1016/j.spasta.2023.100773, 2023.